

# **Crimes and Risks**

by

Jonathan Sarnoff

A dissertation submitted in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy  
(Philosophy)  
in The University of Michigan  
2022

Doctoral Committee:

Professor Gabriel Mendlow, Co-Chair  
Professor Brian Weatherson, Co-Chair  
Professor Sarah Buss  
Professor Scott Hershovitz  
Professor Gideon Yaffe, Yale Law School

Jonathan Sarnoff

[jsarnoff@umich.edu](mailto:jsarnoff@umich.edu)

ORCID iD: [0000-0002-1953-9693](https://orcid.org/0000-0002-1953-9693)

© Jonathan Sarnoff 2022

For my parents.

## ACKNOWLEDGEMENTS

I would first like to thank the members of my doctoral committee for all of the help they have given me, both throughout the process of writing this dissertation and during my graduate education more generally. Brian Weatherson has devoted an unreasonable amount of time to reading (excessively long) drafts of each chapter of this (excessively long) dissertation. His real-time feedback at every stage of my writing has been invaluable, helping me to avoid many (though no doubt not all) of the mistakes in my initial forays into these arguments. Gabe Mendlow first sparked my interest in criminal law theory in a law-school seminar that I took in my first year of graduate school. I am grateful to him for that initial introduction to the subject, which continues to grip me today, and for his guidance as I developed and refined the ideas in this dissertation. Sarah Buss's detailed comments and conversations on the drafts of my chapters have been extremely helpful in clarifying my thinking and steering my revisions. I would also like to thank Scott Hershovitz for participating on my committee as a cognate member. Finally, I would like to thank Gideon Yaffe for his participation as a special member of this committee, especially during the difficult circumstances of this year, which I very much appreciate. He advised the initial law-school paper from which this entire dissertation developed, and I am grateful to him both for first inspiring this project and for helping to shepherd it through the lengthy process it took to reach its present form.

More generally, I would also like to thank everyone else who has been part of the University of Michigan Philosophy Department during the idiosyncratic collection of years when I have been a student here. I cannot imagine a more supportive or vibrant intellectual community of which to have been a member, and I have benefited enormously from its camaraderie during my time here. I am very grateful for the opportunity to have learned with and from so many outstanding

philosophers, and I can only hope in the future to find academic environments that sustain a similarly fruitful and nourishing intellectual community.

An earlier version of chapter one was presented to the University of Michigan Philosophy Department Candidacy Seminar in the fall of 2021, and I am grateful to Adam Waggoner, Alice Kelley, Angela Sun, Arianna Peruzzi, Cameron McCulloch, Gillian Gray, Jamie Tappenden, Laura Soter, Mercy Corredor, Rebecca Harrison, and Sumeet Patwardhan for their discussion and feedback on that draft. An earlier version of chapter three was presented to the University of Michigan Philosophy Department Candidacy Seminar in the fall of 2020, and I am grateful to Alvaro Sottit de Aguinaga, Angela Sun, Arianna Peruzzi, Brendan Mooney, Caroline Perry, Elise Woodard, Jamie Tappenden, Laura Soter, Lingxi Chenyang, Mercy Corredor, Rebecca Harrison, and Sumeet Patwardhan for their feedback and comments on that draft. An even earlier version was presented to the Criminal Law Theory Directed Research seminar at Yale Law School in the fall of 2018, and I am grateful to Amber Qureshi, Charlie Seidell, Emma Stone, Jeremy Lent, Jessica Laird, Kelsey Stimson, Martin Haissiner, Nikila Kaushik, Rachel Rothberg, Retley Locke, Ricky Zacharias, Ryan Martins, Shlomo Klapper, and Tobi Kuehne for their discussion of that draft.

Lastly, I would like to thank Artin Yip, whose steadfast refusal over the past few years to listen to me talk about any of the contents of this dissertation has provided me with a vital anchor to the world beyond, especially when I was most liable to forget it.

# TABLE OF CONTENTS

DEDICATION.....	ii
ACKNOWLEDGEMENTS.....	iii
ABSTRACT.....	vii
Introduction.....	1
Chapter 1: Grading Crime by Its Unintended Results .....	10
I. Criminalization Rules, Grading Rules, and Results.....	23
A. Two Ways to Guide Action .....	26
B. How Criminal Law Guides Conduct.....	34
1. Two Functions of Criminal Law .....	36
2. Criminal Law and Deliberative Criteria .....	40
a. Criminalization Rules .....	40
b. Grading Rules .....	50
C. Guiding Action by Its Results.....	56
1. Criminalization by Results .....	57
a. Excluding Harms .....	57
b. Excluding Risks .....	61
2. Grading by Results .....	71
II. Grading Unintentional Risk-Creation.....	73
A. Results, Proportionality, and Luck.....	75
B. Risk Estimates Through Law.....	84
1. Grading by Risk Factor.....	87
2. Grading Directly by Risk.....	99

C. Risk Estimates by Offenders.....	111
1. Grading by Mental State.....	119
a. Grading by Mens Rea .....	121
b. Grading by Offenders’ Risk Estimates .....	129
2. Grading by Results .....	137
III. Conclusion .....	147
 Chapter 2: Recklessness and the Awareness of Risk.....	 149
I. Reasoning About Risk .....	155
A. Quantitative Reasoning About Risk.....	164
B. Qualitative Reasoning About Risk.....	182
II. Acting Under Uncertainty.....	189
A. Risk and Uncertainty.....	191
B. The Law-Abiding Person.....	195
III. Reckless Mental States .....	209
A. The Nature and Purpose of Action and the Circumstances Known to the Actor .....	211
B. Social Norms and Normative Ignorance.....	223
C. Beliefs About Risk.....	239
IV. Conclusion .....	253
 Chapter 3: Criminal Negligence, Epistemic Recklessness, and the Failure to Inquire .....	 255
I. Negligence as Culpable Ignorance of Risk.....	260
II. Ignorance of Risk and the Failure to Inquire .....	270
III. Criminal Negligence as Epistemic Recklessness.....	281
A. The Culpability of Criminal Negligence.....	282
B. Distinguishing Negligence.....	286
C. The Scope of Criminal Negligence.....	303
D. Tracing .....	310
IV. Conclusion .....	313

## ABSTRACT

This dissertation analyzes three legal doctrines that regulate unintentional aspects of criminal conduct. Chapter one defends the influence the law grants to an action's unintended results in determining the extent of the agent's criminal liability. First, I critique the argument that criminal law's general mens rea requirement allows a result to affect the extent of a defendant's criminal liability only if he possesses mens rea with respect to that result. The rules that define offenses and the rules that specify sentences for individual criminal actions each guide conduct, but they do so differently. Effective offense definitions must be framed in terms of mental states, but effective sentencing rules need not be; therefore, mens rea must set the limits of liability, but the extent of liability need not depend on mens rea alone. Second, I argue that sentencing rules can most effectively guide offenders to reduce the risk that their offenses will unintentionally cause harm if those rules grade the seriousness of each offense based on its results, since the primary alternative, grading offenses based on their risk, requires methods of measuring and communicating risk that exceed the capacities of the criminal justice system.

While chapter one addresses the extent of liability, chapters two and three address its scope: each presents an account of one of the two types of mens rea used to define unintentional crimes. Recklessness, the subject of chapter two, is often analyzed as requiring defendants to possess a belief (or similar mental state) directed towards the degree of risk an action creates. Defining offenses in terms of an agent's beliefs about risk, however, guides individuals to evaluate actions in deliberation by estimating the degree of risk they create. And because individuals ordinarily cannot estimate risks accurately, such rules will not effectively guide them to reduce harm.



Effective rules, I will argue, must identify prohibited actions in a different way—namely, through the features of those actions that make them unreasonably risky. The definition of recklessness does not itself identify those features directly; instead, by requiring individuals to comply with the standard of conduct of the law-abiding person, recklessness indirectly incorporates into law the social norms that regulate dangerous activities. To be reckless, then, defendants must recognize the features of their conduct in virtue of which it violates such a norm.

Finally, chapter three analyzes negligence. Criminal negligence is paradoxical: even though culpability ordinarily requires advertence, negligent defendants are typically understood to be culpable for a risk they inadvertently create. Rather than rejecting culpability for negligence or expanding culpability beyond advertence, I dissolve the paradox by reconceptualizing negligence. Negligent defendants are culpable despite their failure to perceive one risk because they advert to a different risk—the risk that by failing to inquire before acting they will fail to uncover evidence of the risk that they do, in the end, actually fail to perceive. Negligence is epistemic recklessness, the culpable disregard of the risk of inadvertently performing a risky action. Though negligence, like recklessness, involves advertence to risk, it constitutes a distinct form of mens rea because it regulates a different aspect of conduct—the decision to inquire rather than the decision to act. And because omitting inquiry creates only a risk of a risk, negligence is systematically less risky than recklessness, and therefore less culpable.

## Introduction

Each chapter of this dissertation addresses a separate doctrine of criminal law—the unintended results of criminal conduct, recklessness, and negligence—and each offers and defends an analysis that is independent of the arguments advanced in the other chapters. Nonetheless, the three topics resemble one another in important ways: similar divides exist among the arguments scholarly commentators have advanced concerning each doctrine, and the interpretations I propose and defend each stand in similar relationships to the rival views that I will reject. This introduction, then, will discuss the broader theoretical divides that animate disagreement over each particular doctrine, and it will briefly sketch the underlying theoretical unity in the interpretations that I propose. By highlighting that unity, I do not mean to dismiss differences in how my arguments play out in each case; because the three doctrines I analyze are different in many particulars, my theories of each differ in corresponding ways. But the theories of each doctrine that I reject are wrong for broadly similar reasons, in my view, and in each case my proposed alternatives will employ similar strategies for responding to the obstacles I identify, strategies that to some extent reflect a shared conception of what aims criminal law’s doctrines ought to pursue. Thus, I hope to enhance the plausibility of each separate analysis I offer by showing how they all constitute instances of a (hopefully fruitful) general approach to understanding the doctrines of criminal law.

Though the terminology is not without its flaws, many accounts of the doctrines I discuss in this dissertation may be divided between so-called “subjectivist” and “objectivist” approaches.<sup>1</sup> Subjectivist views take a defendant’s

---

<sup>1</sup> For overviews of the distinction, see R.A. Duff, *Subjectivism, Objectivism, and Criminal Attempts*, in HARM AND CULPABILITY 19, 20–21 (A.P. Simester & A.T.H. Smith eds., 1996); and Paul H. Robinson, *The Role of Harm and*

criminal liability to depend on his mental states: as Jeremy Horder explains, the subjectivist view “limits criminal liability to harms or wrongs that are intended or consciously risked.”<sup>2</sup> The law defines intentional crimes, which this dissertation does not address, in terms of the intention to cause a certain harm to perform a certain wrong. Thus, only the consciousness of risk remains to establish criminal liability for unintentional aspects of conduct. Subjectivist approaches accept this position: they require the existence and extent of a defendant’s criminal liability to match the existence and extent of his consciousness of the risk his actions create. The extent of liability for unintended results must depend not on whether those results occur but rather only on the risk that they would;<sup>3</sup> liability for recklessness is justifiable only if a defendant recognized the risk created by his reckless conduct;<sup>4</sup> and negligence, which involves inadvertent risk-creation, is never a justifiable basis for criminal liability.<sup>5</sup>

Each of these judgments concerning the acceptable scope of criminal liability, however, flies in the face both of common practice and of the intuitive judgments that most individuals reflectively endorse about the culpability of particular offenders. Law and morality ordinarily do take the gravity of a wrongful action to depend on whether it causes harm;<sup>6</sup> they do sometimes hold an individual responsible for recklessly disregarding a risk that he does not consciously choose to create;<sup>7</sup> and they

---

*Evil in Criminal Law: A Study in Legislative Deception*, 5 J. CONTEMP. LEGAL ISSUES 299, 299–303 (1994). These labels have widely been criticized as providing an unhelpful and unclear articulation of the difference between these approaches, since subjectivist views take liability to depend on some arguably objective factors, and vice versa. See, e.g., R v. Caldwell [1982] AC 341 (HL) 353 (Lord Diplock) (appeal taken from Eng.) (criticizing “the current vogue for classifying all tests of legal liability as either ‘objective’ or ‘subjective’” because many tests “are not easily assignable to one of those categories rather than the other”).

<sup>2</sup> Jeremy Horder, *A Critique of the Correspondence Principle in Criminal Law*, 1995 CRIM. L. REV. 759, 760; see also Duff, *supra* note 1, at 20 (“We can say that the ‘subjective’ is a matter of the agent’s psychological states . . .”).

<sup>3</sup> E.g., Sanford H. Kadish, *Foreword: The Criminal Law and the Luck of the Draw*, 84 J. CRIM. L. & CRIMINOLOGY 679, 679 (1994) (rejecting doctrines that adjust defendants’ sentences based on whether “the harm they intended or risked does not occur” as “not rationally supportable”).

<sup>4</sup> E.g., LARRY ALEXANDER, KIMBERLY KESSLER FERZAN & STEPHEN J. MORSE, *CRIME AND CULPABILITY: A THEORY OF CRIMINAL LAW* 192 (2009) (“[C]ulpability should attach to those acts of knowingly taking unjustifiable risks . . .”).

<sup>5</sup> E.g., Claire Finkelstein, *Responsibility for Unintended Consequences*, 2 OHIO ST. J. CRIM. L. 579, 581 (2005) (“[N]egligence is incompatible with traditional principles of criminal responsibility.”).

<sup>6</sup> E.g., Kadish, *supra* note 3, at 698 (describing “the popular sentiment that fortuitous results do have a bearing on blameworthiness”).

<sup>7</sup> E.g., R v. Caldwell [1982] AC 341 (HL) 353 (Lord Diplock) (appeal taken from Eng.) (arguing that the “popular or dictionary meaning” of recklessness “covers a whole range of states of mind” beyond consciousness of risk); D.J. Birch, *The Foresight Saga: The Biggest Mistake of All?*, 1988 CRIM. L. REV. 4, 9 (suggesting that with subjectivist accounts of recklessness “the search for a simple formula . . . has triumphed over the need for a

do sometimes take an individual's inadvertence not to excuse him from responsibility but rather to establish his negligence.<sup>8</sup> Objectivist interpretations generally aim to vindicate each of these intuitive judgments. But because the criminal liability they impose cannot be justified by the defendant's recognition of risk, objectivists must identify an alternative basis to justify the imposition of liability. Objectivist theories do not all converge on a single principle of culpability in the same way that subjectivist theories all take a defendant's consciousness of risk to be the appropriate measure of his liability. But objectivism is broadly unified in denying the core subjectivist principle: it insists that the defendant's subjective mental states are not the only legitimate determinants of his criminal liability.<sup>9</sup>

This dispute between the substantive theories of liability defended by subjectivists and objectivists can be understood to reflect differences in the starting points from which each begins its theorizing about unintentional crimes. Subjectivism begins with a plausible general principle of criminal liability, one that ties culpability to an individual's subjective mental states. But applying that principle to evaluate particular defendants produces implausible judgments about their culpability. Objectivism, by contrast, aims to vindicate intuitively compelling assessments of whether and how particular individuals may be held liable. But in doing so it rejects the plausible subjectivist account of what features of an individual's conduct generally justify his liability. Thus, each position in the dispute begins with a thesis plausible when framed at one level of generality, but the two conflict when each is extended more broadly to encompass the subject-matter of the other. Though the particular arguments that subjectivists and objectivists produce differ to some extent between the three subjects addressed in my three chapters, in each case they broadly conform to the pattern I have laid out here. And in each case, I will employ the same general strategy in response. Subjectivism is plausible in its theoretical analysis of the conditions of culpability; objectivism in its particular practical judgments of

---

comprehensive one"); David Ibbetson, *Recklessness Restored*, 63 CAMBRIDGE L.J. 13, 14 (2004) (describing as intuitively reckless a (convicted) defendant who, "in a fit of temper, slammed down a telephone receiver, breaking it . . . [I]t did not occur to him for a moment that he might damage the receiver").

<sup>8</sup> E.g., H.L.A. Hart, *Negligence, Mens Rea, and Criminal Responsibility*, in PUNISHMENT AND RESPONSIBILITY 136, 136 (2nd ed. 2008) ("'I just didn't think' is not in ordinary life, in ordinary circumstances, an excuse . . .").

<sup>9</sup> Duff, *supra* note 1, at 21 ("[Objectivists] deny, rather, that the subjective dimensions of the agent's conduct are all that matter for criminal liability: its 'objective' aspects may also be crucial.").

culpability. But each is typically developed in a manner that conflicts with the other. By contrast, each of my three chapters will propose an account that aims to coherently integrate the compelling aspects of the two rival views: I will attempt to reach objectivist conclusions from subjectivist premises.

Ordinarily, subjectivists defend their claim that an agent's consciousness of risk limits the scope of his criminal liability by appeal to a more general account of the basis of criminal culpability, one rooted in "respect for the moral autonomy of all individuals."<sup>10</sup> On this view, culpability arises from a particular sort of control that individuals exercise over their conduct: a defendant may be culpable only for his improper exercise of rational agency.<sup>11</sup> This approach can easily explain culpability for intentional wrongdoing, since an individual who intends his action to have certain prohibited features exercises rational agency improperly by acting on that intention; instead, individuals must rationally control their behavior by choosing not to perform actions that they recognize to have such features. Obviously, this explanation does not apply directly to individuals who do not intend their actions to have the features because of which they are prohibited. But it can be extended to cover those cases by incorporating the consciousness of risk. Even if an individual does not intend his action to have some feature that the law prohibits, if he is conscious of the risk that it will have that feature he could treat that risk as a reason not to perform it. And by failing to respond appropriately to that reason, he exercises his rational agency improperly. This approach yields the distinctive subjectivist accounts of the three doctrines I analyze. Since the rational force of a risk depends on its magnitude, both the existence and the justifiable extent of a defendant's criminal liability depend on

---

<sup>10</sup> Andrew Ashworth, *A Change of Normative Position: Determining the Contours of Culpability in Criminal Law*, 11 NEW CRIM. L. REV. 232, 235 (2008).

<sup>11</sup> ALEXANDER, FERZAN, & MORSE, *supra* note 4, at 6 ("Such a view presupposes that people act for reasons and that the law can influence those reasons. . . . This approach to preventing harm . . . focuses on the actor's reasons and thus derives its ability to prevent such harms from the capacity and opportunity that agents have to act or abstain from acting for reasons."); Ashworth, *supra* note 10, at 236 ("Underpinning the subjectivist doctrines is a notion of autonomy based on the capacity to reason and to exercise willpower."); Kimberly Kessler Ferzan, *Opaque Recklessness*, 91 J. CRIM. L. & CRIMINOLOGY 597, 641 (2001) (concluding that "consciousness plays a critical role" in avoiding violations of "ought" implies "can"); Finkelstein, *supra* note 5, at 581 ("[O]ur ordinary responsibility practices are predicated on the notion of choice. As such, they extend only to things agents do with awareness of what they are doing or risking."); Kadish, *supra* note 3, at 688 ("[P]unishment is deserved if persons are at fault, and that fault depends on their choice to do the wrongful action, not on what is beyond their control."); Stephen J. Morse, *Reason, Results, and Criminal Responsibility*, 2004 U. ILL. L. REV. 363, 368 (2004) ("Intentional action and forbearance is the only aspect of the human condition that is fully 'up to us,' that is fully within our control, and that can be fully guided by and produced by our reason.").

the degree of risk of which he was conscious. And since negligence involves the failure to advert to a risk, criminal liability for negligence is unjustified.

Thus, subjectivists' approach combines a general analysis of culpability in terms of the improper exercise of rational agency with a focus on one particular kind of reason that might fail to properly influence the exercise of agency—namely, the risks that an agent consciously recognizes an action to create. Objectivist responses typically critique the first of these two elements, suggesting that culpability may instead or in addition depend on factors beyond an agent's subjective awareness of the reasons not to perform an action.<sup>12</sup> By contrast, my arguments will largely accept subjectivists' view that an agent is culpable if he improperly exercises agency by failing to respond appropriately to his awareness of the normatively relevant features of his action or its circumstances. Instead, I will reject the second core subjectivist commitment: in my view, though culpability for unintentional wrongdoing involves a failure to respond appropriately to some normatively relevant feature of an action, that feature need not be the degree of risk that an agent recognizes his action to create. The law imposes liability for a failure to appropriately regulate one's conduct in light of the reasons that ought to bear on one's choices, but individuals do not only—or primarily—regulate their conduct based on the probability that acting will produce a particular legally proscribed harm. Instead, each of my three chapters will argue that criminal culpability can arise from a different kind of improper exercise of rational agency, involving the failure to appropriately respond to a different kind of consideration that constitutes a reason against acting. And this approach, I argue, will vindicate the intuitively compelling judgments that objectivists defend about the influence of results on criminal sentences and about the scope of liability for recklessness and negligence.

In some sense, obviously, the likelihood that an action will produce a legally proscribed harm must be relevant to whether it should or should not be performed: if it matters whether an action would have that feature than it should also matter

---

<sup>12</sup> R.A. DUFF, INTENTION, AGENCY, AND CRIMINAL LIABILITY: PHILOSOPHY OF ACTION AND THE CRIMINAL LAW 129–35 (1990) (discussing the “meanings” of actions); Horder, *supra* note 2, at 763–66 (arguing that performing a wrongful act changes the agent's normative position with respect to consequences that would not be attributable to him based only on how he exercised agency); Kenneth W. Simons, *Punishment and Blame for Culpable Indifference*, 58 INQUIRY 143, 151–56 (2015) (analyzing indifference as a ground of culpability).

whether the action merely might have that feature. But the function of criminal law is not merely to articulate theoretically a list of the considerations that bear on the value of an action; rather, by defining offenses and prescribing punishments for them, criminal liability guides how individuals rationally govern their behavior in practice. And even when an action's likelihood of being harmful may indisputably be relevant to assessing its value in theory, risk nonetheless might not be the most effective consideration for individuals to employ explicitly to govern their conduct. Criminal law does aim to prohibit actions that risk producing harm, but describing the actions prohibited explicitly in terms of risk may not most clearly communicate to individuals the actions they should avoid. And if the law can more effectively guide individuals not to engage in risky conduct by requiring them to employ considerations besides risk itself in their deliberation, an individual may be criminally culpable for failing to respond appropriately to those considerations, rather than only for his failure to appropriately respond to the awareness of risk itself.

Each chapter of this dissertation will broadly pursue this approach in understanding the particular doctrine it analyses. I will approach criminal law as a set of rules designed to rationally guide individuals not to perform certain actions that threaten important social interests. And I will argue that each doctrine I analyze governs an aspect of decision-making that cannot be guided by criminal law most effectively through rules explicitly framed in terms of risk. The rules that grade offenses for the purposes of sentencing guide offenders in choosing among different offenses. Although risk is obviously relevant to how offenders should offend, if they must offend at all, rules that grade by risk will correctly communicate which offenses are more serious only if the adjudicators who apply those rules can correctly measure the risk individual offenses create. I will argue that they cannot reliably do so. Prohibitions on recklessness guide offenders to avoid acts of unjustifiable risk-taking. But if those prohibitions specify unjustifiable risks based on their magnitudes, then individuals must measure the magnitudes of the risks that particular actions create in order to determine whether they exceed the threshold specifying which risks are prohibited. I will argue that they cannot reliably do so. Finally, prohibitions on negligence instruct individuals not to create risks inadvertently. But in their reasoning

individuals cannot directly employ considerations they do not recognize; thus, they cannot be directly guided by the very risks to which they fail to advert. In all three cases, then, rules that explicitly guide individuals' choices according to risk will not effectively lead them to avoid unjustifiably causing harm. Reasoning about risk is hard: we do not always recognize risks, we cannot easily measure them, and we lack a vocabulary for effectively communicating about them. Attempting to devise a strategy for regulating conduct that requires courts and individuals to routinely surmount these difficulties is a loser's game. Criminal law should not play.

If the law cannot effectively guide conduct through rules that explicitly appeal to risk, it should instead devise rules that regulate conduct for the purpose of reducing risk without themselves being explicitly defined in terms of risk. Each chapter of this dissertation will interpret that the doctrine it analyzes as defining such a rule. Grading rules that grade by results communicate only that offenses are more serious if they cause certain results; because offenders themselves can better compare their options in the moment than adjudicators can after the fact, they will better identify which actions are less risky than any grading rules that grade by risk could. Recklessness defines when risk-taking must be avoided not in terms of explicit degrees of risk, which individuals ordinarily cannot judge reliably, but rather in terms of the actual features of actions that make them risky, which individuals can and do recognize accurately. Finally, negligence governs not when risks may inadvertently be created, something not subject to direct rational governance, but rather the care individuals must take in inquiring prior to acting, which enables them to recognize risks and thereby avoid inadvertently creating them.

Failures to govern one's conduct by these rules constitutes the sort of improper exercise of rational agency that gives rise to criminal culpability. And each of these sorts of rules regulate conduct according to some factor other than risk itself. Thus, individuals exercise rational agency improperly not by failing to respond appropriately to the consciousness of risk itself but rather by failing to respond appropriately to the considerations these different rules identify. Since effective grading rules guide individuals in terms of results, not risks, criminal liability must be proportioned to results, not to risk. Since recklessness must define prohibited actions



in terms of their actual features, not in terms of risk itself, whether an individual is reckless depends on what features he recognizes his actions to have, not on how he estimates their probability of causing harm. And since negligence governs inquiry into risk, not the creation of risk, whether an individual is negligent depends on whether he omitted inquiry prior to acting, not on his psychological attitudes towards a risk that he inadvertently creates. In each of these cases, furthermore, the principles that I defend governing criminal liability vindicate the intuitive judgments about the scope and extent of that liability that objectivism defends.

The analytic approach that unites these three chapters, I hope, combines the attractive aspects of both subjectivism and objectivism while avoiding the implausible aspects of each. First, in my view subjectivism correctly emphasizes that criminal law promulgates rules to rationally guide individual conduct, and it correctly concludes that criminal liability must therefore be imposed for individuals' failures to exercise their rational agency in accordance with those rules. But by adopting a more realistic understanding of how rules can guide conduct most effectively in practice, I hope to give a more plausible analysis of which sorts of improper exercise of rational agency justify the imposition of liability. Second, I agree with most objectivist judgments of when and how the imposition of liability is ultimately justified. But objectivists' own defenses of those judgments often focus primarily on their congruence with intuitions; I hope to justify them more persuasively by showing how they are, in fact, consistent with the action-guiding function of criminal law.<sup>13</sup> Furthermore, my focus on how the rules governing the scope and extent of liability guide action enables me to give a more precise account of their exact contours than objectivist accounts that instead defend them by appealing to often vague and imprecise moral intuitions.<sup>14</sup>

---

<sup>13</sup> Michael Moore defends grading by results, for example, in large part by denying that reasons must be given to actually support grading by results—either because the relevance of results is self-evident, *see* Michael S. Moore, *The Independent Moral Significance of Wrongdoing*, 5 J. CONTEMP. LEGAL ISSUES 237, 259–63 (1994), or because the intuitive judgments in particular cases that results are relevant themselves constitute premises in an argument for the importance of results, *see id.* at 263–71. Certainly, people who can't think of reasons in support of their positions respond quite pragmatically by instead giving reasons why they shouldn't be expected to give reasons in support of their positions. But this approach seems unlikely to be particularly successful at persuading those who do not already accept the views that one argues one need not argue in favor of.

<sup>14</sup> Many objectivists, for example, argue that recklessness and negligence depend on a defendant's culpable indifference—his morally blameworthy failure to be appropriately motivated by a concern for the interests of others. *E.g.*, DUFF, *supra* note 12, at 161–67. But how much is one supposed to be motivated by those concerns? Which interests is one obligated to consider? And which particular actions fail to display the appropriate level of concern? This approach is entirely silent on these questions, and thus provides virtually no guidance to juries or

Obviously, each of the proposals advanced in this dissertation will depend on its strength, taken individually. But I hope that taken together they can demonstrate the viability of the approach they share, which analyzes of criminal law's doctrines focuses squarely on the pragmatic question of how the prohibitions it enacts can effectively guide individuals' deliberation and choice.

---

individuals themselves as to which acts are actually reckless or negligent—it would hardly be less helpful for the law to say merely that no act is guilty without a guilty mind. For a related criticism, see John Gardner & Heike Jung, *Making Sense of Mens Rea: Antony Duff's Account*, 11 OXFORD J. LEGAL STUD. 559, 578 (1991).

## **Chapter 1:**

# **Grading Crime by Its Unintended Results**

Criminal law often defines criminal conduct in terms of its effects on the world. Although the law characterizes prohibited actions in part by identifying the mental states that an offender must possess to be guilty of performing them, it also characterizes those actions in part by describing their results on people or things—injuring or killing another person, say, or destroying another’s property. Since such definitions allow the results an action causes to determine whether it satisfies the definition of a crime, they allow criminal liability to depend on what results an action actually produced, not merely on whether it was intended or expected to produce them. But an action’s results depend on a host of factors, of which the nature of the defendant’s own conduct is only one, though often an important one. Because these other factors can affect an action’s results, a defendant’s criminal liability for offenses defined in terms of results may not depend on the nature of his conduct alone. Offenders with the same intentions and beliefs who perform the same bodily movements involving the same objects may nonetheless differ in their criminal liability if, because of fortuitous differences in their circumstances, their actions produce different results. An offender’s sentence may depend on whether the means he chose to perform the crime fails to produce the result that he intended—due, say, to the flight path of a bullet—or on whether additional results such as injury or even death occur that he did not intend—due, say, to a victim’s poor health.<sup>1</sup> In these

---

<sup>1</sup> *E.g.*, MODEL PENAL CODE § 5.05 (1) (AM. LAW INST. 1962) (providing that the attempt to commit a first-degree felony is a second-degree felony); CAL. PENAL CODE § 664 (providing that the attempt to commit a crime should be punished by an indeterminate prison term with a maximum length of half the sentence for the completed crime); PA. CONS. STAT. ANN. 18 § 2502(b) (providing that a person commits murder by causing death in the course of performing another felony).

cases, a difference in an action's results can on its own change the agent's criminal liability without any corresponding difference in the agent's intentions or beliefs in acting.

Thus, although criminal law widely employs results to define criminal conduct,<sup>2</sup> the practice has been criticized frequently.<sup>3</sup> To be sure, nobody denies that the law's reliance upon results accurately reproduces a central aspect of how we evaluate human conduct in general: we do typically judge actions in part by the results they actually produce.<sup>4</sup> But while critics concede that results may intuitively be relevant to evaluating conduct, they deny that any rational basis can be identified to justify their use in determining individuals' criminal liability. According to these arguments, criminal law simply has no reason whatsoever to evaluate offenses differently if they differ in their results without any difference in the agent's intentions, beliefs, or conduct: in Sanford Kadish's words, the use of results to determine criminal liability is "not rationally supportable" in that it "does not serve the crime preventive purposes of the criminal law, and is not redeemed by any defensible normative principle."<sup>5</sup> For these critics, the dispute over whether results should affect criminal liability does not juxtapose compelling arguments that must be

---

<sup>2</sup> Sanford H. Kadish, *Foreword: The Criminal Law and the Luck of the Draw*, 84 J. CRIM. L. & CRIMINOLOGY 679, 679 (1994) (describing the practice's "near universal acceptance in Western law"); George P. Fletcher, *What Is Punishment Imposed for?*, 5 J. CONTEMP. LEGAL ISSUES 101, 106 (1994) (noting that "[v]irtually everybody insists" that "wrongdoing is aggravated by the occurrence of harm").

<sup>3</sup> See LARRY ALEXANDER, KIMBERLY KESSLER FERZAN & STEPHEN J. MORSE, *CRIME AND CULPABILITY: A THEORY OF CRIMINAL LAW* 171–96 (2009); ADAM SMITH, *THE THEORY OF MORAL SENTIMENTS* 117–19 (Knud Haakonssen ed., 2002); Andrew Ashworth, *Belief, Intent, and Criminal Liability*, in *OXFORD ESSAYS IN JURISPRUDENCE (THIRD SERIES)* 1 (John Eekelaar & John Bell eds., 1987); Andrew Ashworth, *Taking the Consequences*, in *ACTION AND VALUE IN CRIMINAL LAW* 107 (Stephen Shute, John Gardner, & Jeremy Horder eds., 1993); Larry Alexander, *Crime and Culpability*, 5 J. CONTEMP. LEGAL ISSUES 1–30 (1994); Joel Feinberg, *Equal Punishment for Failed Attempts: Some Bad but Instructive Arguments Against It*, 37 ARIZ. L. REV. 117 (1995); Kadish, *supra* note 2; Kimberly D. Kessler, *The Role of Luck in the Criminal Law*, 142 U. PA. L. REV. 2183 (1994); Stephen J. Morse, *Reason, Results, and Criminal Responsibility*, 2004 U. ILL. L. REV. 363 (2004); Stephen J. Schulhofer, *Harm and Punishment: A Critique of Emphasis on the Results of Conduct in the Criminal Law*, 122 U. PA. L. REV. 1497 (1974); J.C. Smith, *The Element of Chance in Criminal Liability*, 1971 CRIM. L. REV. 63 (1971).

<sup>4</sup> E.g., ALEXANDER, FERZAN, & MORSE, *supra* note 3, at 174–75; SMITH, *supra* note 3, at 118; Kadish, *supra* note 2, at 688–89.

<sup>5</sup> Kadish, *supra* note 2, at 680; see also ALEXANDER, FERZAN, & MORSE, *supra* note 3, at 174–75 (analogizing the question of whether results should influence liability to whether it matters "if one's victim has eleven toes, or the killing occurs on a Tuesday, or the shooter uses his left hand, or the reckless driver's car is red, and so forth"); Bjorn Burkhardt, *Is There a Rational Justification for Punishing an Accomplished Crime More Severely than an Attempted Crime*, 1986 BYU L. REV. 553, 571 (1986) ("Is there a rational justification for punishing a completed crime more severely than an attempted crime? I am not certain that the reasons which are given for a less severe punishment are 'rational.'"); Feinberg, *supra* note 3, at 118 (describing differences based on results as "more arbitrary than rational"); Schulhofer, *supra* note 3, at 1599 (asking "whether there is any perspective . . . no matter how controversial or idiosyncratic, from which distinctions based on harm can be justified").

carefully balanced to determine which ultimately is more convincing; instead, opponents insist forcefully that treating results as relevant to liability is simply irrational.<sup>6</sup> And given the absence of any possible rational justification on one side of the dispute, they understand it instead as a conflict between the rationality of law and criminal justice and the non-rational, emotional force of our instinctive responses to others' wrongdoing.<sup>7</sup>

To be sure, it may not be obvious how that conflict should be resolved. One strategy for defending the use of results is not to provide reasons for why the results of an action should on their own affect the agent's criminal liability but rather to deny that the inability to provide such reasons shows the use of results to be unjustified. On this view, our intuitive responses to others' wrongdoing require no explicit rational justification to be vindicated: we may accept our instinctive moral commitments even if we cannot produce reasons to defend them. Different kinds of arguments have been given for why ordinary moral emotions may dictate the law's rules for determining criminal liability even absent any rational justification for those rules. Some such arguments appeal to moral considerations. Perhaps we intuitively find bad results morally significant because they just are worse intrinsically: no reason can be given for why because the importance of results is simply a brute fact, a principle of morality that does not rest on any more fundamental principle and thus cannot be justified by appeal to one.<sup>8</sup> Or perhaps, since punishment expresses the community's

---

<sup>6</sup> E.g., ALEXANDER, FERZAN, & MORSE, *supra* note 3, at 174 (noting "the contrast between the strength of the choice argument and the weakness of the results argument"); Andrew Ashworth, *A Change of Normative Position: Determining the Contours of Culpability in Criminal Law*, 11 NEW CRIM. L. REV. 232, 255 (2008) ("Why should the intentional or knowing commission of a crime of the same family render a person criminally liable for the unforeseen consequences of the conduct . . . ? No good reason has been given for this, other than the assertion that there has been a change of normative position, and that begs the question."); Feinberg, *supra* note 3, at 122 ("What puzzles me is how bad many of [the arguments for using results] are.")

<sup>7</sup> ALEXANDER, FERZAN, & MORSE, *supra* note 3, at 174 ("The only argument on the table for punishing results more . . . is a phenomenological argument. We feel guiltier and blame others more when harm occurs. . . . [N]one of us can find a compelling principle beyond these purported emotions in the case of results."); SMITH, *supra* note 3, at 118 (describing the importance intuitively according to results as "an irregularity in the sentiments of all men"); Feinberg, *supra* note 3, at 119 ("Arbitrariness is the absence of rule, as in the bare will of an authority who can exert his power free of accountability, in a manner without rhyme or reason, which in turn makes predictability and security from abuse difficult, and fairness an inapplicable notion . . ."); Fletcher, *supra* note 2, at 107 ("This ongoing dispute presents, in the end, a conflict between two cultures in criminal law, one tied to common sense, the other reflecting what Bruce Ackerman once dubbed 'scientific policy making.'"); Kadish, *supra* note 2, at 699 ("If I am right that some of the law's doctrines are not rationally defensible . . . it is our moral judgments that are irregular.")

<sup>8</sup> See Michael S. Moore, *The Independent Moral Significance of Wrongdoing*, 5 J. CONTEMP. LEGAL ISSUES 237 (1994).

moral attitudes towards offenders, the influence of results within morality requires that they influence punishment, as well.<sup>9</sup> Other arguments appeal to political considerations. Perhaps democratic legitimacy requires that the law not deviate too sharply from the moral intuitions of those it governs.<sup>10</sup> Or, perhaps a legal system like ours simply cannot function properly, given the role it assigns to juries, if ordinary citizens reject its principles.<sup>11</sup> These approaches defend the law's use of results without even attempting to explain why the results of an action should make a difference to criminal liability.

In this chapter, I will attempt to provide an alternative kind of defense for the role results play in determining criminal liability—in particular, an account that explains why results should make a difference to an individual's criminal liability. Rather than conceding that no reasons can justify the influence criminal law accords to results, then defending that influence even in the absence of any such reasons, I will directly confront the charge of irrationality levied so commonly against their use.<sup>12</sup> Thus, I will argue that there do exist sound reasons why the results of an action

---

<sup>9</sup> For an approach broadly along these lines, see R.A. DUFF, *INTENTION, AGENCY, AND CRIMINAL LIABILITY: PHILOSOPHY OF ACTION AND THE CRIMINAL LAW* (1990) (“The answer must be that the fact of her failure should qualify these responses to her because it matters: it matters to us, and should matter to her, whether her criminal attempt succeeded or failed, just because it matters to us [and should matter to her] whether or not her victim actually suffered the material harm she tried to cause him.”); and R.A. Duff, *Auctions, Lotteries, and the Punishment of Attempts*, 9 *LAW & PHIL.* 1, 36 (1990) (“[W]e are, and he should be, relieved that it failed . . . ; and if his punishment is to communicate to him an adequate understanding of his crime, it should surely aim to communicate the appropriateness of this relief.”).

<sup>10</sup> Feinberg, *supra* note 3, at 125 (“In a democracy, Holmes insisted, those who make, enforce, and interpret the law cannot move too far out of harmony with the common people, else the law will seem tyrannical at worst, obscure at best.”); Fletcher, *supra* note 2, at 107 (“If punishment is tied to shared cultural judgments of wrongdoing, then harming is clearly a greater wrong than trying or preparing to harm. It follows that the result, when it occurs, is part of the action for which an offender is punished.”); Kadish, *supra* note 2, at 702 (“There are limits, therefore, particularly in a democratic community like ours, to how far the law can or should be bent by reformers to express a moral outlook different from that of the deeply held intuitive perceptions of the great mass of humanity, irrational though they may seem to some.”).

<sup>11</sup> *E.g.*, Schulhofer, *supra* note 3, at 1511 (“[D]eference to desires for vengeance, on the part of victims or the public generally, is necessary to promote some other value.”). The authors of the Model Penal Code were particularly concerned that juries would engage in nullification were the sentences imposed for particular crimes to deviate too greatly from how jurors assessed their gravity, as might be the case were the results of crimes not to affect the sentences imposed for performing them. MODEL PENAL CODE § 2.03 cmt. 1 (AM. LAW INST. 1985) (“How far the penal law ought to attribute importance in the grading of offenses to the actual result of conduct, as opposed to results attempted or threatened, presents a significant and difficult issue. Distinctions of this sort are essential, at least when severe sanctions are involved, for it cannot be expected that jurors will lightly return verdicts leading to severe sentences in the absence of the resentment aroused by the infliction of serious injuries. Whatever abstract logic may suggest, a prudent legislator cannot disregard these facts in the enactment of a penal code.”)

<sup>12</sup> Of course, by providing an argument that rationalizes the use of results to determine criminal liability, I do not mean to reject other arguments that defend the use of results without purporting to explain why they should affect criminal liability. It may be true both that reasons do exist justifying the use of results to determine liability and that the intuitive relevance of results could justify their use even without those reasons.

should at least sometimes affect the agent's criminal liability. As I have noted, critics typically reject the use of results in unusually uncompromising terms, not merely endorsing arguments against it but rather characterizing it as an irrational<sup>13</sup> abandonment of scientific policymaking.<sup>14</sup> This chapter, then, will refute that critique by identifying reasons why the dependence of criminal liability on results does, in fact, advance the purposes of the law. The importance of results is not merely an irrational quirk of our moral emotions; from the standpoint of scientific policymaking, too, results sometimes should matter to criminal liability. One function central to criminal law is to promulgate rules of conduct to guide individuals' behavior: by prohibiting and punishing certain actions, it requires individuals not to perform actions of that type. The criminal law must therefore impose liability to effectively communicate the rules individuals must employ to govern their conduct. And, in my view, communicating those rules effectively requires criminal liability sometimes to depend on the results of conduct.<sup>15</sup> The widespread dismissal of the use of results as purely irrational is thus mistaken: at least sometimes an action's results should affect the agent's criminal liability for performing it.

Many scholars discuss the influence of results on criminal liability as if it were a single, unified phenomenon that must be defended in all instances, if at all.<sup>16</sup> But there are many different ways in which an action's results might affect the agent's liability for performing it, and in my view there is no reason to think that all

---

<sup>13</sup> *E.g.*, Kadish, *supra* note 2, at 679.

<sup>14</sup> Fletcher, *supra* note 2, at 107.

<sup>15</sup> Many commentators have focused on the criminal law's role in guiding behavior in order to criticize its use of results: since individuals cannot control whether an action actually does produce its expected result, they argue, the law cannot effectively guide individuals' behavior by requiring them actually to produce certain results. *See, e.g., id.* (attributing opposition to the use of results to "the theoretical conviction that control over one's bodily movements determine the outer limit of culpability and liability"); Kessler, *supra* note 3, at 2194 ("The actor has very little control over whether her actions will ultimately end in their intended result."); Morse, *supra* note 3, at 365 ("[U]nderstanding the centrality of the action-guiding function clarifies and perhaps solves many traditional puzzles of criminal liability, such as the relation of results to desert . . ."). Of course, I agree that individuals may often be mistaken about the actual results of their actions, and thus that it is not always possible to control what results one's action produces. I differ from critics of the use of results, however, in that I will undertake a more detailed analysis of exactly how the criminal law guides behavior through its imposition of punishment. Certainly, individuals do lack full control over what results their actions produce, but I will argue that the use of results to determine criminal liability best guides individuals' choices concerning aspects of their conduct that they can control.

<sup>16</sup> Alexander, *supra* note 3, at 1 ("[M]y topic in this paper . . . is whether the central organizing principle of criminal liability should be the social harm caused by defendant's [sic] act or should instead be the defendant's culpability as reflected in his act."); Fletcher, *supra* note 2, at 106 ("Yet the question remains whether wrongdoing is aggravated by the occurrence of harm."); Kessler, *supra* note 3, at 2237 ("Although we cannot altogether avoid the role of chance in our lives, it should not be a part of the criminal law.").

such proposals must stand or fall together. Thus, in two respects I will more narrowly defend only a particular use of results in criminal law. First, at least two different kinds of results might affect criminal liability: intended results and unintended results. Whether criminal liability should depend on the occurrence of an action's intended results is simply the problem of criminal attempts, an "ancient problem" at which "[e]very bona fide philosopher of law tries his hand at least once."<sup>17</sup> But I will not consider intended results in this chapter. Instead, I will focus on the many other doctrines that allow the unintended results of an action to affect the agent's criminal liability. For example, perhaps the most notorious example in contemporary American law of the effect of results on criminal liability, the doctrine of felony murder, concerns killings caused unintentionally during the commission of other felonies.<sup>18</sup> Thus, I will argue only that an individual's criminal liability should sometimes depend on the unintended results of his action; I will not consider how, if at all, my arguments could be extended to the problem of criminal attempts.

Furthermore, my argument for allowing unintended results to affect criminal liability depends on the particular relevance of risk to unintended results. The guidance the law provides individuals concerning whether they may intentionally cause prohibited harms is quite simple—they may not. Because all actions that intentionally cause harm share a particular mental state—namely, the intention to cause the harm—the law can guide behavior in terms of that mental state by forbidding individuals from acting on that intention. But no particular mental state distinguishes all actions that cause unintended harms, for such actions bear only a causal relation to the harm they produce. Thus, the law cannot guide individuals to avoid such harms by requiring them not to act with some particular mental state, for no such mental state is present in all actions that unintentionally cause harm. Instead, the law can protect against unintended harms by guiding individuals to reduce the risk

---

<sup>17</sup> Feinberg, *supra* note 3, at 117. Other attempts at the problem include R.A. DUFF, *CRIMINAL ATTEMPTS* (1996); GIDEON YAFFE, *ATTEMPTS: IN THE PHILOSOPHY OF ACTION AND THE CRIMINAL LAW* 310–33 (2010); Andrew Ashworth, *Criminal Attempts and the Role of Resulting Harm under the Code, and in the Common Law*, 19 *RUTGERS L.J.* 725 (1988); Lawrence C. Becker, *Criminal Attempt and the Theory of the Law of Crimes*, 3 *PHIL. & PUB. AFFAIRS* 262 (1974); Omri Ben-Shahar & Alon Harel, *The Economics of the Law of Criminal Attempts: A Victim-Centered Perspective*, 145 *U. PA. L. REV.* 299 (1996); Duff, *supra* note 9; Jerome Hall, *Criminal Attempt — A Study of Foundations of Criminal Liability*, 49 *YALE L.J.* 789 (1940); and David Lewis, *The Punishment that Leaves Something to Chance*, 18 *PHIL. & PUB. AFFAIRS* 53 (1989).

<sup>18</sup> See, e.g., N.Y. PENAL LAW § 125.25(3) (2019).



that their actions will create such a harm. That rule, I will argue, is sometimes communicated most effectively by allowing criminal liability to depend on results.

Just as two different kinds of results might affect criminal liability, those results might affect criminal liability in one of two different ways. First, the law might define some crimes wholly in terms of results: whether a particular action satisfied that definition and thus fell within the scope of the criminal prohibition would depend only on whether it caused a particular result, regardless of any of its other features. Alternatively, results might determine only how the law grades the seriousness of a particular crime: out of all actions that satisfy the definition of a crime, the law might impose a higher sentence on offenders who cause a particular result. In practice, criminal liability depends on results only in grading; almost never do results alone suffice for an action to satisfy the definition of a crime.<sup>19</sup> Crime generally requires mens rea: an individual commits a crime not merely by causing a result but rather by causing it with a particular mental state.<sup>20</sup> Thus, the minimum conditions of criminal liability cannot be met simply by causing a particular result; instead, the result must be caused with a particular mental state to be a crime.

But while no action lacking mens rea can be a crime, no mens rea requirement applies to doctrines that merely distinguish among crimes based on their seriousness: one crime may be more serious and carry a higher sentence than another even without any corresponding difference in mens rea. Thus, the results of a crime may on their own increase how the law judges its seriousness, and how harshly it will sentence the offender. Felony murder, for example, employs results alone to determine which defendants are guilty of murder: those who cause death during the commission of a felony are guilty, whereas those who do not cause death are not.<sup>21</sup> But because the doctrine applies only to deaths caused during the commission of another felony,

---

<sup>19</sup> See Paul H. Robinson, *The Role of Harm and Evil in Criminal Law: A Study in Legislative Deception*, 5 J. CONTEMP. LEGAL ISSUES 299, 303–04 (1994) (“The subjectivist view seems to be the overwhelming modern view of the minimum requirements for criminal liability.”).

<sup>20</sup> There are some exceptions to the requirement of mens rea, and thus sometimes causing a particular result may suffice on its own for criminal liability. In particular, many regulatory schemes are enforced through public welfare offenses, which prohibit causing a violation of those regulations regardless of whether mens rea exists. See Darryl K. Brown, *Public Welfare Offenses*, in *THE OXFORD HANDBOOK OF CRIMINAL LAW* 862 (Markus D. Dubber & Tatjana Hörnle eds., 2014). Some other crimes are defined not to require any mens rea with respect to circumstances: prohibitions on statutory rape that require only that the victim be under the age specified by statute are perhaps the most notorious example. See, e.g., *R v. Prince* [1875] 2 L.R.C.C.R. 154 (Eng.).

<sup>21</sup> See N.Y. PENAL LAW § 125.25(3) (2019).

results alone do not determine whether any action is a crime at all: any act to which the felony murder doctrine applies would still constitute the underlying felony, even were it not to cause death. Thus, felony murder does not determine which actions are crimes but rather determines only how serious of a crime a particular action is—whether it is murder or some lesser felony—and thus what sentence the offender will receive. In general, results typically affect criminal liability only in this way: the results that an offense causes can affect how serious of an offense it is and thus how severe of a sentence it carries.<sup>22</sup> Consequently, I will defend the use of results only to affect how crimes are sentenced. Grading crime by its unintended results, I will argue, communicates most effectively that offenders must reduce the risk their offenses create.

This chapter will have two parts. For opponents of grading by results, the omission of a mens rea requirement from grading rules, which govern the sentencing of particular criminal acts, is inconsistent with the inclusion of such a requirement in criminalization rules, which specify the minimum requirements of criminal liability. Mens rea is a condition for criminal liability, in their view, because it is necessary for culpability, but if culpability requires mens rea in general then some difference in mens rea must be required for some difference in culpability to justify sentencing one offender more harshly than another.<sup>23</sup> The first part of my argument will consider and reject this objection. Criminalization rules and grading rules guide conduct in different ways, and mens rea is important only to the former kind of guidance. Thus, though mens rea is required for any action to be a crime, consistency does not demand that it be required whenever an offender's sentence is increased. Instead, mens rea is correctly employed only by those rules that cannot guide conduct effectively without it.

These two kinds of rules employed by criminal law guide conduct differently, in my view, by providing different kinds of criteria for individuals to employ in deliberation and choice to evaluate possible actions. Some criteria simply identify a

---

<sup>22</sup> Robinson, *supra* note 19, at 300 (“[M]ost American jurisdictions take a subjectivist view of the minimum requirements of liability but many simultaneously take an objectivist view of grading.”).

<sup>23</sup> *E.g.*, Kenneth W. Simons, *When is Strict Criminal Liability Just?*, 87 J. CRIM. L. & CRIMINOLOGY 1075, 1105 (1997) (“For it is implausible that retributive theory would condemn strict liability in criminalizing without paying any heed to strict liability in grading.”).

dimension along which some actions are better or worse than others. This sort of guidance does not provide any conclusive determination as to whether an agent may choose any particular action; instead, it identifies one consideration relevant to choice but defers to the agent's own judgment as to how to balance it against any others that may apply. Other criteria, however, purport to decide themselves whether an agent may perform a particular action: rather than allowing agents to determine whether the balance of reasons supports an action, such criteria conclusively exclude certain options as ineligible to be chosen. Criminalization rules promulgate the latter kind of criterion: agents are required to exclude actions that fall within the definition of a crime. The fact that an action is prohibited by law is not merely one consideration to be weighed against others in deliberation; rather, criminally prohibiting an action conclusively determines that it may not be chosen. Grading rules promulgate the former kind of criterion. By identifying some offenses as more serious than others, they guide agents to prefer less serious offenses in deliberation, rather than more serious ones. But grading rules do not identify which actions must be excluded from deliberation entirely. Indeed, since grading rules grade only crimes, which must all be excluded from deliberation, no further guidance of that sort remains to provide. Instead, by distinguishing better and worse ways of violating criminalization rules, grading rules guide the deliberation only of individuals who are willing to disregard the conclusive guidance criminalization rules provide.

Differences between these kinds of guidance impose different constraints on how each kind of rule must be structured. Criminalization rules guide conduct by identifying options that agents must exclude from deliberation, while leaving them free to choose among the options that remain. I will argue that criminalization rules employing results alone would fail to guide conduct effectively, then, because they fail to distinguish effectively between these two kinds of option. Defining an offense in terms of results alone criminalizes any action that causes that result, regardless of how unlikely it was. Because criminalization rules identify the actions that agents must exclude from deliberation, such rules would require agents to exclude an action from deliberation on any occasion in which it actually would cause an unlikely harm. But in deliberation agents can employ only the information they actually possess:

they can decide whether or not a particular action must be excluded from deliberation only on the basis of the information they possess about that action. Thus, if agents possess identical information about two actions, they cannot treat those actions differently in deliberation, even if only one of the two would actually cause harm; instead, they must treat identically all options about which they possess identical information, by excluding either all or none. And because any action has some possibility of causing a particular harm, any option available to an agent is indistinguishable from some action that actually would cause harm. Criminalization rules that employ results alone, then, provide no basis for distinguishing between the options that must be excluded from deliberation and those that are eligible to be chosen.

To distinguish effectively between those categories of actions, a criminalization rule must instead define prohibited actions in terms of information about them that the agent possesses and that, therefore, she may employ in deliberation. Criminalization rules that employ mens rea satisfy this requirement by defining offenses in terms of the conscious mental states with which they are performed, such as the intention to perform a certain type of action or the awareness that an action would be of that type.<sup>24</sup> Since those mental states have as their content information about particular actions, agents may determine which options to exclude from deliberation on the basis of their information about their available actions: they must exclude an action if, in light of the information they possess about it, performing it would be a crime. And because the options about which individuals possess no such information are therefore eligible to be performed, criminalization rules framed in terms of mens rea do distinguish which available options must be excluded from deliberation, as is required for them to guide conduct effectively.

Grading rules guide conduct differently than criminalization rules do: by identifying some offenses as more serious than others, they guide offenders to perform offenses that are less serious. Crucially, this guidance does not require

---

<sup>24</sup> Because negligence involves the failure to perceive a risk rather than the conscious disregard of one, it is less clear that offenses of negligence are identified in terms of the offender's conscious mental states. Compare MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962) with *id.* § 2.02(2)(d). In chapter three of this dissertation, though, I will defend an analysis of negligence that does interpret it in terms of the offender's conscious mental states.

classifying all actions available to an agent into two categories. To guide offenders to perform offenses that are less serious, grading rules need only identify some basis along which offenders may compare different offenses, but so long as offenders can compare offenses along that dimension there is no reason to define some privileged threshold to divide them into two categories. A rule framed in terms of results alone faces no difficulty providing this sort of guidance: if causing a particular result is bad, then actions likelier to cause it are worse than actions less likely to cause it.

Identifying some actions as better or worse than others would not suffice for the sort of guidance criminalization rules aim to provide: the fact that one action is worse than another does not, on its own, entail whether an agent may perform both, neither, or only one. But grading rules do not need to provide guidance concerning that further step: they guide conduct simply by identifying which offenses are better or worse. Unlike criminalization rules, grading rules can guide conduct effectively without requiring mens rea; thus, the omission of a mens rea requirement from grading rules is justified even though the mens rea is required in criminalization rules.

This argument, of course, merely rejects an objection against grading by results. This chapter will go on, then, to present a positive case for grading by results. My argument will primarily be comparative: grading by results is justified, in my view, because it is better than the alternatives. I have suggested that one function of grading rules is to guide offenders to perform offenses that are less likely to cause certain harmful results; thus, I will focus on how effectively different methods of grading offenses can discharge this function.<sup>25</sup> The primary alternatives I will consider are various forms of grading by risk—grading rules that use some assessment of the risk created by a particular offense in order to determine the sentence to be imposed for performing it. Certainly, if such rules could be implemented successfully, they could effectively guide conduct, as well: imposing more severe sentences on offenses that create greater risks would guide offenders instead to choose offenses that create lesser risks. But successfully implementing such rules would require some methodology for accurately estimating the risks created by

---

<sup>25</sup> The risk of causing harm is only one factor that makes some offenses more serious than others. Thus, while some grading rules aim to reduce that risk, others guide offenders to perform offenses that are less serious in other ways.

individual actions. And, I will argue, any plausible method of estimating risk will be accurate only to a very limited extent. Consequently, grading by risk can effectively play only a limited role in guiding offenders to perform less risky offenses. Furthermore, the inaccuracy of grading by risk undermines arguments in its favor even outside of the particular function I ascribe to grading rules. Grading inaccurately by risk is simply not grading by risk at all. Thus, the inaccuracy of all plausible methods of estimating risk undermines any case for grading by risk regardless of what reasons are proposed in its favor.

The most obvious way to estimate risk for the purpose of sentencing would be for some government entity to estimate it. Two steps are generally involved in estimating risk. First, the underlying information about each offense to be used in those estimates must be identified. The government entities that could most plausibly perform this task are the courts that sentence offenders. Second, that information must actually be used to estimate risk. More variation exists in which entity could perform this task. Courts themselves could directly estimate risk in addition to identifying the information used to produce those estimates. But rules determining how to estimate risk could also be produced in advance: lawmakers could write grading rules that individually specify how each aspect of an offense would affect its risk. Courts, on this approach, would first determine which risk factors characterized a particular offense, then impose the sentence required by all the grading rules that apply. Nonetheless, I will argue, on any of these strategies the entities that estimate risk will do so inaccurately, introducing errors through each of these two steps required to produce risk estimates. And since grading rules guide conduct through how they impose punishment, these inaccuracies will undermine the effectiveness of that guidance: grading by risk will guide offenders to perform the offenses that the law inaccurately judges to be less risky, not those that actually are less risky.

If government entities are likely to estimate risk inaccurately, whose risk estimates should guide offenders in deciding which offenses to perform? There is an obvious alternative: offenders themselves. How, then, could a grading rule guide offenders to perform actions that they themselves judge to be less risky? One natural suggestion would be to grade offenses by offenders' own risk judgments: punishing

offenders more severely for performing offenses that they themselves judge to be riskier would guide them to perform offenses that they judge to be less risky. Just as criminalization rules determine guilt based on an offender's mental states, then, grading rules might themselves employ mental states to grade offenses. Grading by mental state, however, cannot guide conduct effectively. Because risk estimates directed towards risks one disregards have no natural connection either to sensory inputs or behavioral outputs, they are extremely difficult for courts to identify. Furthermore, while employing offenders' own risk estimates to determine sentences might be effective when offenders do actually estimate risks, there is no reason to think they will always do so. And if sentences can be increased only by risks that offenders consciously recognize and estimate, grading by mental state will simply guide offenders to avoid estimating risks, not to avoid creating them.

Instead, to guide offenders to perform the offenses that they themselves judge least risky, grading rules should grade offenses by their results. By punishing offenses more severely if they cause certain results, grading rules simply guide offenders to perform offenses that will not cause those results. Various forms of grading by risk all produce some model for estimating risk, then instruct individuals to reduce risk as estimated by those models. Such rules thus guide offenders to perform offenses that are less risky with respect only to those features whose effect on risk is captured in the law's models. By contrast, under grading by result sentencing depends not on whether some model recognizes an aspect of an offense as affecting its risk but rather on whether performing that offense will cause a particular harm. Thus, to avoid offenses that will be punished more severely offenders should simply strive to avoid causing harm. To be sure, given the uncertainty we face about the results of our actions, even offenders who aim to avoid causing harm sometimes will cause it. But that uncertainty is simply unavoidable: given limited human knowledge it is not possible to identify which offenses will or will not cause harm based only on the information accessible to offenders at the time they perform their offenses. And the best way to avoid harm, given that uncertainty, is for offenders themselves to act based on their own best efforts at identifying how to avoid it, as grading by results guides them to do.

## I. Criminalization Rules, Grading Rules, and Results

Results might play one of at least two different roles within the criminal law: whether an action causes a particular result might determine whether or not that action is a crime at all, or it might determine only the seriousness of an action that already qualifies as a crime. In practice, as I have noted, the law employs results almost exclusively (though not entirely) for the latter purpose: while it is common for crimes to be graded by their results, it is rare for an action's results alone to determine whether it is a crime at all, independent of the agent's intentions, knowledge, purposes or other mental states.<sup>26</sup> Thus, there is a difference between how the law sets the minimum conditions of criminal liability and how it calibrates the sentences imposed on offenders: while the results of an action may on their own affect the sentence imposed for performing it, individuals ordinarily must possess some kind of mens rea with respect to a harmful result for their action to be a crime.

Many have criticized this difference in the use of mens rea in criminalization and in grading. Criminalization rules require an offender to possess mens rea to be guilty of a crime, they argue, because the offender's culpable mental state justifies his criminal liability.<sup>27</sup> As the maxim puts it, *actus non facit reum nisi mens sit rea*: no act is guilty without a guilty mind.<sup>28</sup> And if a culpable mental state is required to justify imposing criminal liability at all, it might seem that the extent of an offender's justifiable criminal liability must be determined by the degree of culpability of his mental state.<sup>29</sup> On this approach, there should be no difference between the

---

<sup>26</sup> Robinson, *supra* note 19, at 303–04.

<sup>27</sup> *E.g.*, Ashworth, *supra* note 3, at 117 (attributing “[t]he preoccupation with the mental or fault element in serious criminal offenses” to “the central significance of culpability in the criminal law”); Smith, *supra* note 3, at 64 (“As society became more enlightened, the law began to insist upon a requirement of fault. It must be proved that the accused intended the harm, or caused it recklessly . . . or at least negligently . . . . So liability came to be based, not solely on the harm done, but on harm done in a blameworthy manner.”).

<sup>28</sup> *E.g.*, 3 EDWARD COKE, INSTITUTES OF THE LAWS OF ENGLAND 10 (London, E. and R. Brooke 1797). As John Gardner and Jeremy Horder have both noted, however, taken literally the maxim requires only a guilty mind, not a mind guilty with respect to every feature determining the seriousness of an offense. John Gardner, *Rationality and the Rule of Law in Offences Against the Person*, 53 CAMBRIDGE L.J. 502, 509–510 (1994); Jeremy Horder, *A Critique of the Correspondence Principle in Criminal Law*, 1995 CRIM. L. REV. 759, 770. Thus, on its face the requirement of mens rea appears to apply only to questions of criminalization, not to questions of grading.

<sup>29</sup> *E.g.*, Ashworth, *supra* note 3, at 116 (arguing that the censure expressed by criminal punishment must be restricted both “to those who deserve it” and “to the extent that they deserve it”); Kadish, *supra* note 2, at 688 (“Isn’t desert the same whether or not the harm occurs? It is commonly accepted that punishment is deserved if persons are at fault, and that fault depends on their choice to do the wrongful action, not on what is beyond their control.”).



requirements for finding an individual criminally liable at all and the requirements for increasing his criminal liability.<sup>30</sup> Since an offender must possess some mens rea for the imposition of any criminal liability to be justified, he must similarly possess some additional mens rea for the imposition of any additional criminal liability to be justified.<sup>31</sup> Finding an offender who lacks a more serious kind of mens rea to be guilty of a more serious crime would be a form of strict liability no less than finding guilty an offender lacking any mens rea at all.<sup>32</sup> And since most theorists reject strict liability in general, they reject strict liability in grading, too.<sup>33</sup>

Doctrines that impose a more severe sentence on offenders who cause certain results regardless of whether they possess any form of mens rea towards those results obviously do not require the extent of an offender's criminal liability to match the degree of culpability of his mental states. Consequently, those who insist that mens rea must apply both to criminalization and to grading typically reject grading rules that grade by results alone.<sup>34</sup> Instead, such theorists propose revisions in the structure

---

<sup>30</sup> As Paul Robinson notes, the dispute over the criminal law's use of results has often been framed as a dispute between unified "objectivist" and "subjectivist" accounts of culpability that each address both criminalization and grading, thereby obscuring even the possibility that differences in the functions of criminalization and grading rules might justify differences in the roles that results play in each. Robinson, *supra* note 19, at 303–04.

<sup>31</sup> Alexander, *supra* note 3, at 4 ("If social harm is not necessary for criminal desert—and no one believes it is sufficient—its alleged ability to enhance criminal desert becomes mysterious."); Ashworth, *supra* note 6, at 255–56 ("What is required is a reason for saying that, given that the basis for liability is the subjective and cognitivist requirement of intentional conduct, the extent of the liability should be governed not by what the actor intended or foresaw (as subjectivists, applying the correspondence principle, would insist) but rather by events outside the actor's contemplation."); Simons, *supra* note 23, at 1121 ("[A]dding a penalty to the felony because it resulted in death seems no more justifiable than punishing someone for an accidental, non-negligent homicide today simply because he committed a felony last year.").

<sup>32</sup> See, e.g., Stuart P. Green, *Six Senses of Strict Liability: A Plea for Formalism*, in APPRAISING STRICT LIABILITY 1, 3 (A.P. Simester ed., 2005) (defining strict liability to encompass offenses that omit a mens rea requirement with respect to only some elements, even if an actions that failed to satisfy that element at all would still constitute a crime by satisfying all other elements); Simons, *supra* note 23, at 1082–85 (distinguishing between strict liability with respect to criminalization and strict liability with respect to grading). Although my argument in this chapter could be described as a defense of strict liability in grading, I will avoid that terminology, since enough inconsistency exists in the use of the term "strict liability" that employing it is likelier to create more confusion than it avoids. See Green, *supra*, at 2–9 (listing six different meanings with which the term is used).

<sup>33</sup> E.g., Simons, *supra* note 23, at 1105 ("For it is implausible that retributive theory would condemn strict liability in criminalizing without paying any heed to strict liability in grading."). But see Kenneth W. Simons, *Is Strict Criminal Liability in the Grading of Offences Consistent with Retributive Desert?*, 32 OXFORD J. LEGAL STUD. 445 (2012) (defending a limited role for strict liability in grading).

<sup>34</sup> Ashworth, *supra* note 3, at 124 (criticizing grading by results for "neglect[ing] the pivotal importance of culpability in both liability and sentencing"); Kessler, *supra* note 3, at 2194–95 ("Hence the causal/result requirement conflicts with our view that rational agents should be held accountable for their actions: it allows fortuity to determine their crime. This is something upon which criminal responsibility should not depend."); Morse, *supra* note 3, at 366 ("I conclude that the consequences of action cannot be fully guided and are thus not appropriate predicates for desert. Full culpability and desert are established by intentional action that risks a harmful result."); Smith, *supra* note 3, at 64 ("It is the importance still attached by the law to the requirement of harm that leads to most of the problems that I am going to discuss.").

of criminal law that would prevent results alone from affecting the scope of offenders' criminal liability, such as a requirement that some form of mens rea correspond to every element of each offense,<sup>35</sup> offense definitions framed solely in terms of the defendant's conduct,<sup>36</sup> or explicit grading rules that determine sentences using only the defendant's mental states.<sup>37</sup>

In this Part, I will provide a general argument for why the requirement of mens rea in criminalization rules need not imply that mens rea is required to justify any increase in an offender's sentence. In my view, a culpable mental state plays a necessary role in setting the minimum conditions of criminal liability that it need not play in determining the seriousness of—and thus the sentence imposed for—performing a particular criminal act.<sup>38</sup> My argument will depend on an analysis of exactly how the rules imposed by criminal law guide the conduct of the individuals they bind. In particular, the rules that define which kinds of actions are crimes and the rules that identify the seriousness of each particular criminal act guide individuals' conduct in different ways. The kind of guidance a rule aims to provide, furthermore, constrains what form that rule may take. In particular, while the guidance that criminalization rules provide requires offenses to be defined in terms of the mental states with which they are performed, no such requirement applies to the kind of guidance that grading rules provide. Consequently, results may legitimately play a different role in each kind of rule. I will begin, then, by explaining two different ways in which rules of conduct might guide individuals' behavior, which correspond to the kinds of guidance provided by the two kinds of criminal rule I will discuss. This account of how rules may guide conduct, in turn, will identify what conditions a rule must satisfy if it is to guide effectively—and, in particular, will identify how those

---

<sup>35</sup> This is the “correspondence principle” defended by some scholars of English criminal law. *See, e.g.*, Ashworth, *supra* note 6, at 245 (requiring “correspondence between the level of harm D chose to cause or risk causing and the level of harm for which D is held criminally liable”); Andrew Ashworth & Kenneth Campbell, *Recklessness in Assault—And in General?*, 107 *LAW Q. REV.* 187, 192 (1991) (“In general the criminal law seems to adopt what may be termed a principle of correspondence between mens rea and actus reus: if the offense is defined in terms of certain consequences and certain circumstances, the mental element ought to correspond with that by referring to those consequences or circumstances.”).

<sup>36</sup> Feinberg, *supra* note 3, at 119 (“I would prefer to have a more comprehensive crime with a new technical sounding name, like ‘Wrongful Homicidal Behavior.’”).

<sup>37</sup> ALEXANDER, FERZAN, & MORSE, *supra* note 3, at 282–83.

<sup>38</sup> For other arguments for this position, see Mitchell N. Berman, *Blameworthiness, Desert, and Luck*, *NOÛS*, 16–18 (forthcoming); John Gardner, *The Wrongdoing that Gets Results*, 18 *PHIL. PERSP.* 53, 61–63 (2004); Horder, *supra* note 28; and Michael Otsuka, *Moral Luck: Optional, not Brute*, 23 *PHIL. PERSP.* 373, 374–82 (2009).

conditions differ between the kinds of guidance that criminalization rules and grading rules aim to provide.

*A. Two Ways to Guide Action*

In general, individuals decide how to act by selecting one course of action from a set of available options. Deliberation begins with an agent who must choose which action to perform, and it ends when the agent somehow identifies which of his options to choose. In order to evaluate those options and select among them, agents must employ some criteria that together single out the particular action that will be performed. The choice of which criteria to employ in selecting an option, in turn, substantially influences which option an agent will ultimately choose, since by employing different criteria he is likely to settle on different actions. Consequently, one way to influence individuals' behavior is to influence the criteria that they employ when selecting which action to perform. What kinds of criteria might influence individuals' behavior in this way? Broadly speaking, individuals employ two kinds of criteria in deliberation and choice, each of which might be employed to guide their behavior.<sup>39</sup>

First, some criteria give agents a reason of some particular strength that weighs for or against some or all of the options that are available. For example, the fact that a particular action might benefit someone is a consideration that weighs in favor of performing it, and the fact that something is expensive is a consideration that weighs against buying it. Thus, individuals may be guided by instructing them to prefer, say, more beneficial actions over less beneficial ones: an agent who employs that criterion in deciding what to do will value certain actions higher or lower than she otherwise would have, and in some cases this difference in evaluation may change what action she ultimately chooses. Rules requiring individuals to employ a criterion of this sort influence their conduct by making them likelier to choose an action that the criterion evaluates relatively positively and less likely to choose one that it evaluates relatively negatively.

---

<sup>39</sup> I do not mean to claim, obviously, that this is the only kind of division that might be drawn among the principles that individuals employ in practical reasoning.

But though criteria of this sort can influence behavior, they do so by identifying only one consideration that is relevant to an agent's decision; crucially, they do not fully dictate what conclusion an agent must reach about what to do. Identifying actions as better or worse than others along a specified dimension does not, on its own, instruct the agent how to translate that judgment into action, for whether or not to perform a particular action ordinarily depends on the full range of reasons that weigh for or against it, not merely on any one such reason taken individually. The fact that an action might benefit another, say, does genuinely constitute a reason to perform it, but that fact alone does not conclusively settle whether to perform it, since even a harmful action may be justified in some contexts, given sufficiently strong countervailing considerations. Criteria of this sort, then, influence an agent's decision but do not fully determine it: they identify a consideration that should influence an agent's decision without dictating the choice an agent ultimately must make. Rather, agents themselves must determine which action is best justified by the balance of reasons. Thus, though a criterion of this sort does influence an agent's decision, it does not displace the agent's own judgment, since the agent must still herself evaluate what option the balance of reasons supports once this particular criterion is included.

By contrast, the second sort of criterion an agent might employ in selecting which action to perform does take this further step. That is, it does not merely identify one reason among many that weighs for or against different options; instead, it explicitly tells agents which options they may choose and which they may not. A criterion of this sort, for example, might exclude acts of torture, or might rule out eating meat but not eating fish, or might prohibit the use of racial slurs. Unlike criteria that merely identify one consideration that agents must consider in deciding what to do, criteria of this sort are decisive—if they apply to an action they do conclusively settle whether or not it may be performed. If racial slurs are prohibited, then the fact a sentence uses such a slur conclusively settles whether one may utter it: one simply may not. To be sure, exceptions plausibly exist to many decisive criteria, though the existence and scope of such exceptions is often disputed. The prohibition on killing may have exceptions in cases of self-defense or wartime; others—say, the

prohibition on rape—may have no exceptions whatsoever. Thus, applying a particular criterion may require both determining how the criterion applies and determining whether any exceptions apply, as well. (There may even be exceptions to the exception, and so forth.) Furthermore, different criteria may be more or less difficult to apply—it is plausibly easier to tell which actions are excluded by a requirement to avoid racial slurs, say, than by a requirement to speak respectfully. But though different decisive criteria do differ in this way, all are alike in the kind of influence they exert on individuals' decision-making: ascertaining whether or not the criterion excludes a particular action can itself dictate the conclusion to deliberation, since the criterion determines conclusively whether an action is eligible to be performed, regardless of any other considerations.

Consequently, the manner in which an individual must employ a decisive criterion in practical reasoning differs from the manner in which she must employ a non-decisive one. Because the latter sort of criterion merely identifies one consideration relevant to which actions agents should choose, they cannot employ a non-decisive criterion alone to decide how to act; instead, though all non-decisive considerations influence an agent's decision, the ultimate choice depends on the cumulative effects of those considerations together, not on any one individually. By contrast, the influence of a decisive criterion does not depend in this way on the potentially countervailing influence of other applicable criteria. If a decisive criterion applies to a particular option, that in itself settles how the agent must respond—namely, by avoiding that option and excluding it as an eligible choice in her deliberation. If the use of racial slurs is prohibited, one may not choose a sentence that involves a slur (at least when no exception applies, should any exist). The options to which a decisive criterion applies (and to which no exception applies) are ineligible to be chosen, regardless of what other considerations may weigh in their favor: the criterion says simply not to choose them. To guide individuals' conduct by requiring them to employ a decisive criterion of this sort, then, influences their decisions not merely by instructing them to include a particular consideration as an element in their reasoning; rather, it purports to conclusively determine what conclusion that reasoning must reach, at least partially—it conclusively determines what choice

individuals must make with respect to certain options, though it may leave open multiple options between which individuals are free to choose. Such criteria do displace an agent's own judgment, at least partially: they do not allow agents themselves to evaluate the balance of reasons that apply to certain options but rather require the agent to defer to the judgment that those options must be excluded.<sup>40</sup>

Obviously, it is an interesting philosophical question how an account of practical reasoning can explain the difference between these two distinct ways in which a consideration might influence an individual's decision-making.<sup>41</sup> Nonetheless, my arguments in this chapter do not depend on how the theory of practical reasoning accommodates the distinction between decisive and non-decisive criteria. Rather, however the distinction may ultimately be explained, my concern here is with its implications for rules that guide individuals' behavior. Rules can influence behavior by providing individuals with a criterion to employ in practical reasoning. Since there are two kinds of criteria that individuals might employ in their practical reasoning, there exist two kinds of corresponding rules that guide individual behavior by providing individuals with a particular criterion to employ—namely, rules that guide behavior according to either a decisive or a non-decisive criterion. And differences in how decisive and non-decisive criteria influence practical

---

<sup>40</sup> I am certainly not the first to draw this sort of distinction between two ways in which a consideration may enter into an individual's practical reasoning. Most notably, Joseph Raz has similarly distinguished between different kinds of reasons on the basis of how conflicts between them are resolved. JOSEPH RAZ, PRACTICAL REASON AND NORMS 35–36 (1999). Conflicts between some kinds of reasons, which Raz terms first-order reasons, are resolved based directly on the weight of reasons on each side of the conflict, so that “one ought, all things considered, to do whatever one ought to do on the balance of reasons.” *Id.* at 36. Such reasons correspond to what I have called non-decisive criteria. Another kind of reason, which Raz calls a second-order reason, functions very differently—it resolves conflicts between reasons not by affecting which side the balance of reasons favors but rather by directly dictating a resolution regardless of how the balance of reasons weighs. *See id.* at 35–39. These reasons are what I have called decisive criteria. For a survey of theoretical issues in the analysis of how reasons conflict with one another, see generally Errol Lord & Barry Maguire, *An Opinionated Guide to the Weight of Reasons*, in WEIGHING REASONS 3 (Errol Lord & Barry Maguire eds., 2016).

<sup>41</sup> In Raz's view, these kinds of reason influence decision-making in different ways because they are reasons for different kinds of things. First-order reasons, as the name indicates, are reasons for or against a particular action, but second-order reasons are instead reasons for or against acting on a particular first-order reason. RAZ, *supra* note 40 at 39–40. This analysis explains why decisive criteria on their own can conclusively determine how an agent must decide—namely, because such reasons themselves require the agent to ignore any countervailing reasons. Nonetheless, I do not mean to endorse (or to reject) Raz's account of the difference between what I have called decisive and non-decisive criteria; I mention his analysis merely to note one way to explain the two different sorts of practical reasoning that I have described here. Others exist as well; for example, Ruth Chang argues that decisive criteria are not reasons concerning which reasons an individual should act upon but rather are reasons concerning which kind of choice situation an agent should place herself in. *See* Ruth Chang, *Comparativism: The Grounds of Rational Choice*, in WEIGHING REASONS 213, 220–27 (Errol Lord & Barry Maguire eds., 2016).

reasoning, I will argue, require corresponding differences in the rules that impose each sort of criterion on individual decision-making.

Decisive criteria purport to exclude certain options entirely from individuals' deliberations about what to do: they identify which options available to an individual are ineligible to be chosen, thereby restricting deliberation and choice to the option or options that remain. The choice of which actions to include as options in deliberation and which to exclude is a binary one: either an action will be included or it will not be. No third category exists containing options that are neither included in deliberation nor not included in deliberation. A rule can provide individuals with a decisive criterion to employ in practical reasoning, then, only if it identifies, for every option available, which category it belongs in—whether it is ineligible for choice and therefore must be excluded from deliberation, or whether it is available for deliberation and choice. Such rules might give explicit conditions specifying when an action falls into each category, but they need not. More naturally, they might instead identify only those actions that are (or are not) excluded, implying that all remaining actions fall into the remaining category; alternatively, they might classify actions into more than two categories, so long as each category is itself identified as containing excluded or eligible actions. Because either of these approaches classifies all options into one of two categories, it identifies in some way which of those options must be excluded from deliberation.

But a rule that fails to identify which actions are excluded and which are eligible cannot provide individuals with a decisive criterion to employ in their deliberation. An agent must treat each option as eligible or as excluded—any given option will either be included in deliberation or it will not be—and a decisive criterion guides that choice by instructing an agent which options to exclude. If a rule leaves an option unclassified between those two categories, then, it simply fails to give agents the kind of guidance that decisive criteria purport to provide: it does not tell agents whether to include or exclude that option in their deliberation. The rule may communicate that information in various ways, but it must communicate it somehow. Agents cannot be instructed which options to exclude unless a rule somehow divides options into the excluded and the eligible: a rule that fails to

classify some or all options into one of those two categories fails to instruct agents either to exclude those options or not to exclude them. Instructing individuals that it is better to drive more slowly, for example, fails to provide them with a decisive criterion: because it does not classify actions into two categories but rather ranks them along some scale, it does not on its own identify any actions that must be excluded from deliberation. Which actions would those be? How fast is too fast to drive? The rule does not say. Similarly, no decisive criterion is provided by the rule that walking is more environmentally friendly than taking public transit, which in turn is more environmentally friendly than driving. Though such a rule does group actions into categories and compare those categories, it fails to identify which categories, if any, must be excluded from deliberation. Only the worst? All but the best? The rule does not say. By contrast, a rule forbidding speeds above one hundred miles per hour does impose a decisive criterion on individuals' decision-making: it instructs agents to exclude from deliberation all options involving higher speeds, and it leaves agents free to select from among only those options that remain.

Because a rule providing individuals with a non-decisive criterion to employ in their decision-making guides deliberation in a different manner, it need not provide the same kind of information to communicate the criterion individuals must follow. Whether to include an option in deliberation is a binary choice, for each option must either be excluded from deliberation or be treated as eligible for choice. The binary nature of this choice, in turn, requires a rule with a binary structure, which classifies each option into one of two categories. But a non-decisive criterion does not aim to dictate the result of any binary choice to individuals—it does not influence decision-making by dividing options into two groups that receive different treatment in an agent's deliberations. Rather, a non-decisive criterion influences decisions by identifying some respect (not necessarily involving two distinct categories) in which certain actions are better or worse than others. Depending on what other options are available and what other reasons apply, this additional consideration may change which action an agent judges to be best supported by the balance of reasons, since an option that would otherwise be suboptimal might instead be chosen in light of the additional reason identified in the rule.



A non-decisive criterion guides behavior, then, by guiding how agents make relative comparisons among their choices. Ascertaining which option from some set is favored by the balance of reasons is essentially a comparative task: agents deliberate over which action to perform by comparing available options along various normative dimensions and ultimately selecting the one that is best overall. Thus, the choice of a particular option to perform reflects only the comparative judgment that the option selected is better than any of the available alternatives. A non-decisive criterion influences behavior by influencing how agents judge the balance of reasons when comparing those options during deliberation: sometimes, an agent who would otherwise judge one option superior to another will make the opposite judgment, and hence choose a different action, once the additional criterion identified by some rule is also incorporated into deliberation. That is, a non-decisive criterion can influence behavior if the overall balance of reasons between options changes once it incorporates the sort of comparison between those options that the criterion requires individuals to make. Non-decisive criteria do not influence choice by requiring agents to take a particular attitude towards individual options—say, excluding them from deliberation entirely. Rather, by identifying one particular basis of comparison that agents must consider in their deliberations, they affect an agent’s comparative judgments about which options are better supported than others by all applicable reasons.

Consequently, a rule requiring individuals to employ a particular a non-decisive criterion needs only to communicate comparative information about groups of options—in particular, it must inform agents how to compare options on one particular basis.<sup>42</sup> With this information, agents can employ the criterion in deliberation to identify which option is better supported by the balance of reasons: an agent who would otherwise have judged one option better than another by a weak

---

<sup>42</sup> Such rules will influence deliberation more substantially to the extent that they are more precise about how the criterion they identify must be weighed against other criteria in evaluating the balance of reasons—that is, about the exact margin by which different actions are better or worse than others. Nonetheless, it is still possible to influence deliberation even without providing information about weight: such a rule may influence agents who did not recognize that a particular consideration was normatively significant at all, for example, or who incorrectly took a consideration to weigh against an action when it in fact weighs in favor. (Alternatively, these examples might show that information about weight is necessary but can be very vague: to specify only that a consideration weighs in favor of an action specifies that it must be given positive weight without specifying the magnitude of that weight.)

enough margin may, after applying the criterion, reach a different conclusion about what action to perform. Normally, a rule will convey that comparative information by specifying that actions possessing some feature to a greater extent are better or worse by some magnitude than actions possessing it to a lesser extent (or not at all). Such a rule specifies that actions possessing that feature must sometimes be chosen over actions lacking it even when an agent would otherwise judge the latter action superior. So long as the margin by which the agent judges one action superior is smaller than the margin by which the criterion imposed by the rule favors the other, the criterion will shift the balance of reasons.

Thus, although rules may influence individuals' behavior by requiring them to employ either a decisive or a non-decisive criterion, differences in the deliberative role of the two kinds of criteria require the rules imposing each to differ structurally. As I have argued, the former sort of rule must make a particular kind of comparison between an agent's options: because it aims to guide individuals in making the binary choice of whether to exclude actions from deliberation, it must classify every option into one of two categories. That classification decisively settles an agent's decision with respect to at least some available options by excluding them entirely, while leaving the agent free to follow her own judgment with respect to the actions that remain. That binary classification is central, because the rule provides guidance solely through how it classifies actions: actions excluded by the rule may not be chosen, but the rule does not aim to influence how an agent chooses among actions in each category. By contrast, a non-decisive criterion merely compares options along a particular dimension without making conclusive judgments about whether any particular action may be performed; the considerations it identifies are merely inputs to deliberation, and they may always be outweighed given sufficiently strong countervailing reasons. Such rules merely identify one way in which some actions are preferable to others while leaving open the conclusion that an agent should reach in any particular case. Consequently, such rules need only compare options in some respect; they need not make a binary classification of options.<sup>43</sup>

---

<sup>43</sup> Other legal theorists have also introduced and discussed a distinction broadly along these lines, but they have applied it primarily in contexts outside of criminal law. Ronald Dworkin's account of adjudication, for example, relies heavily on a distinction between two kinds of considerations that courts may employ to reach a decision,

Rules must satisfy fewer restrictions, then, in order to impose non-decisive criteria on individual decision-making than to impose decisive criteria: certain rules that cannot impose a decisive criterion on individuals can still impose a non-decisive one. The rule that slower driving is to be preferred because of safety, as I noted, does not identify any decisive criterion, because it fails to specify which actions must be excluded from deliberation. Nonetheless, because a non-decisive criterion needs to identify only some basis for comparing options, regardless of whether that comparison involve a binary classification, that rule does convey a non-decisive criterion—namely, one that favors slower over faster driving. A rule that cannot communicate a decisive criterion may still communicate a non-decisive one: more rules can impose non-decisive criteria on individual deliberation than can impose decisive criteria, and rules that purport to provide decisive criteria for individuals to employ must satisfy certain conditions, to be effective, that do not apply to rules purporting to provide non-decisive criteria.

### *B. How Criminal Law Guides Conduct*

Many kinds of rules purport to guide individual behavior—rules of morality, of custom, of etiquette, and so forth. Law is one such set of rules. Indeed, the law’s role in guiding conduct is so central that it forms the basis of many jurisprudential accounts of law. A prominent account of the rule of law, for example, holds that legal rules must be capable of guiding the conduct of those they govern: as Joseph Raz puts it, the “the basic intuition from which the doctrine of the rule of law derives” is the

---

which Dworkin calls rules and principles. Ronald M. Dworkin, *The Model of Rules*, 35 U. CHI. L. REV. 14, 22–29 (1967). Like rules that impose what I have called decisive criteria, rules “are applicable in an all-or-nothing fashion. If the facts the rule stipulates are given . . . the answer it supplies must be accepted” (if it is valid). *Id.* at 25. By contrast, a principle, like a rule that imposes what I have called non-decisive criteria, “states a reason that argues in one direction, but does not necessitate a particular decision.” *Id.* at 26. Robert Alexy draws a similar distinction in similar terms in developing his theory of constitutional adjudication. ROBERT ALEXY, *A THEORY OF CONSTITUTIONAL RIGHTS* 47–48 (Julian Rivers tran., 2010). Rules articulate exactly what they forbid; they constitute “*fixed points* in the field of the factually or legally possible.” *Id.* at 48. Thus, like decisive criteria, rules on their own determine what agents must do. Principles instead merely “require that something be realized to the greatest extent possible given the legal and factual possibilities.” *Id.* at 47. Thus, what a principle requires in a particular case depends on other factors—what Alexy terms the legal and factual possibilities—in the same way that the ultimate application of non-decisive criterion depends on what other reasons weigh for or against a particular choice. Both Dworkin and Alexy develop this distinction to analyze how judges decide cases, not to decide how ordinary individuals make choices. But as Raz notes, there is no good reason to restrict the application of the distinction in this way. Joseph Raz, *Legal Principles and the Limits of Law*, 81 YALE L.J. 823 (1972) (“Professor Dworkin creates the impression that the norm-subjects of all legal principles are the courts. But other people and institutions may also be the norm-subjects of principles . . .”).

claim that “the law must be capable of guiding the behavior of its subjects.”<sup>44</sup> This criterion, Raz argues, must be met for law to effectively achieve any of its aims, just as “[a] knife is not a knife unless it has some ability to cut.”<sup>45</sup> Criminal law, then, must aim to guide the conduct of those it governs. Indeed, the law itself acknowledges this objective: the first item the Model Penal Code lists among the “general purposes of the provisions governing the definition of offenses” is “to forbid and prevent conduct that unjustifiably and inexcusably inflicts or threatens substantial harm to individual or public interests.”<sup>46</sup> I have argued, though, that there are two ways in which legal (or other) rules are capable of guiding conduct, which differ in the kinds of criterion they require the individuals they govern to employ in their deliberation. And these differences determine, in part, what the structure of that rule must be. Which of these two kinds of guidance does the criminal law aim to provide?

This question, in my view, need not receive the same answer across criminal law as a whole. Instead, different elements of criminal law differ in whether they influence conduct by requiring individuals to employ a decisive or a non-decisive criterion in their deliberation. The body of legal doctrines that constitute the criminal law incorporate at least two different kinds of rule, which I will call criminalization rules and grading rules.<sup>47</sup> These two kinds of rule differ in the legal function they serve in the adjudication of criminal cases: criminalization rules define offenses by describing when a defendant’s conduct will constitute a particular crime, while grading rules specify how serious a given offense is for the purpose of determining what sentence will be imposed on the offender. Furthermore, I will argue, these two kinds of rules differ in which kind of criterion they aim to impose on individuals’ deliberation: criminalization rules impose decisive criteria on deliberation whereas grading rules impose non-decisive criteria.

---

<sup>44</sup> Joseph Raz, *The Rule of Law and its Virtue*, in *THE AUTHORITY OF LAW: ESSAYS ON LAW AND MORALITY* 210, 214 (1979); see also LON L. FULLER, *THE MORALITY OF LAW* 70 n.29 (1964).

<sup>45</sup> Raz, *supra* note 44, at 226; FULLER, *supra* note 44, at 97 (“[W]e are concerned, not with the substantive aims of legal rules, but with the ways in which a system of rules for governing human conduct must be constructed and administered if it is to be efficacious and at the same time remain what it purports to be.”).

<sup>46</sup> MODEL PENAL CODE § 1.02(a) (AM. LAW INST. 1962).

<sup>47</sup> Of course, these are not the only kinds of rules found in the criminal law, which at the very least also includes rules governing excuses and rules adjusting sentences based on factors other than the elements of the offense of conviction, such as criminal history or post-conviction behavior. No doubt there are other sorts of rules, as well.

### *1. Two Functions of Criminal Law*

The distinction between criminalization rules and grading rules is displayed most clearly by provisions of criminal law that explicitly separate the criminalization rule and the grading rule that correspond to a particular offense. The Model Penal Code section covering forgery is representative.<sup>48</sup> Its first clause, entitled “Definition,” states what conduct constitutes forgery:

A person is guilty of forgery if, with purpose to defraud or injure anyone, or with knowledge that he is facilitating a fraud or injury to be perpetrated by anyone, the actor:

(a) alters any writing of another without his authority; or

(b) makes, completes, executes, authenticates, issues or transfers any writing so that it purports to be the act of another who did not authorize that act, or to have been executed at a time or place or in a numbered sequence other than was in fact the case, or to be a copy of an original when no such original existed; or

(c) utters any writing which he knows to be forged in a manner specified in paragraphs (a) or (b).

“Writing” includes printing or any other method of recording information, money, coins, tokens, stamps, seals, credit cards, badges, trade-marks, and other symbols of value, right, privilege, or identification.<sup>49</sup>

This clause states a criminalization rule. It specifies what a defendant must do to commit the crime of forgery: he must alter or make a writing (a term that is itself defined) in a particular way, or utter a writing altered or made in that way, with a particular purpose or particular knowledge. Thus, the rule defines a certain kind of behavior as criminal. This is the characteristic function of criminalization rules: they specify what behavior constitutes a crime.

The forgery provision’s second clause, “Grading,” classifies different kinds of forgery as felonies of different degree:

Forgery is a felony of the second degree if the writing is or purports to be part of an issue of money, securities, postage or revenue stamps, or other instruments issued by the government, or part of an issue of

---

<sup>48</sup> MODEL PENAL CODE § 224.1.

<sup>49</sup> *Id.* § 224.1(1).

stock, bonds or other instruments representing interests in or claims against any property or enterprise. Forgery is a felony of the third degree if the writing is or purports to be a will, deed, contract, release, commercial instrument, or other document evidencing, creating, transferring, altering, terminating, or otherwise affecting legal relations. Otherwise forgery is a misdemeanor.<sup>50</sup>

This is a grading rule. It specifies how serious of an offense different kinds of forgeries are —namely, that forgeries are most serious if they involve government instruments or property claims, less serious if they involve other sorts of legal documents, and least serious if they involve any other sort of writing. As this provision illustrates, grading rules employ a scale along which they grade the seriousness of different offenses. The Model Penal Code, for example, divides crimes into felonies, misdemeanors, and petty misdemeanors, in order of decreasing seriousness;<sup>51</sup> felonies are themselves divided into those of the first, second, and third degree, again in order of decreasing seriousness.<sup>52</sup> These evaluations of seriousness, in turn, are employed primarily in sentencing:<sup>53</sup> the seriousness of a particular crime partially determines what sentence will be imposed on a defendant found to be guilty of committing it.<sup>54</sup> This is the characteristic function of grading rules: they specify the seriousness of different kinds of criminal conduct in order to assign different sentences to individual criminal acts.

Criminalization rules and grading rules, then, jointly determine how a defendant’s conduct influences his criminal liability: the criminalization rule establishes whether that conduct was a crime, and the grading rule partially specifies how it will be sentenced. Together, these rules govern a criminal adjudication: a jury or judge will employ the first to determine whether a defendant committed a crime, then the second to grade that crime for the purposes of sentencing. (Other factors may

---

<sup>50</sup> *Id.* § 224.1(2).

<sup>51</sup> *Id.* § 1.04(1)–(4).

<sup>52</sup> *Id.* § 6.01(1).

<sup>53</sup> *E.g., id.* § 6.01(1) (“Felonies defined by this Code are classified, *for the purpose of sentence*, into three degrees, as follows.” (emphasis added)).

<sup>54</sup> *Id.* §§ 6.05–.09. As indicated both by the indeterminacy of sentences authorized by the Model Penal Code and by the range of sentences authorized for each degree of felony, many factors besides the degree of the offense committed may also influence a defendant’s sentence. Those factors may include characteristics of the defendant’s criminal conduct that the code does not incorporate into its grading rules, as well as factors other than offense characteristics.

also affect the sentence imposed or served.) By defining and grading each offense, Model Penal Code provisions ordinarily promulgate both a criminalization rule and a grading rule, though not all criminalization and grading rules are as neatly separated into different clauses as in the forgery provision.<sup>55</sup> Other systems of criminal law likewise promulgate both criminalization rules and grading rules, though they may differ substantially from the Model Penal Code in how they structure the legal provisions that promulgate them.

Rather than including one rule of each kind in a single provision, for example, federal criminal law instead largely separates the two kinds of rule between two entirely separate sources of law. Federal criminalization rules are found in federal statutes codified primarily in Title 18 of the U.S. Code. Sections of Title 18 define various criminal offenses. The forgery and counterfeiting provisions, for example, are found in chapter 25: section 471 defines the crime of forging or counterfeiting obligations or securities of the United States,<sup>56</sup> section 503 the crime of forging or counterfeiting a postmarking stamp,<sup>57</sup> and so forth. These provisions promulgate criminalization rules: they specify that certain conduct is a crime. In addition, they include some grading rules—namely, they specify the maximum (and sometimes minimum) sentence that may be imposed for each offense.<sup>58</sup> But the vast majority of federal criminal law’s grading rules are part of a wholly separate body of law—the Sentencing Guidelines promulgated by the United States Sentencing Commission.<sup>59</sup>

---

<sup>55</sup> To be sure, not every offense definition promulgates a criminalization rule. The code defines some offenses in terms of others: a person commits robbery, for example, if “in the course of committing a theft” he seriously injures another, threatens to do so, or commits another first- or second-degree felony. *Id.* § 222.1. From the standpoint of criminalization, this offense definition is entirely redundant: because robbery is simply theft committed in a particular manner, all robberies already fall within the definition of theft. Defining the crime of robbery, then, does not criminalize any conduct that is not already criminalized. Robbery is a particular kind of aggravated theft, and the definition of robbery could just as well have been included among the provisions grading thefts: they define when a theft is a third-degree felony, misdemeanor, or petty misdemeanor, and robbery defines when a theft is a first- or second-degree felony. *Compare id.* § 223.1(2) *with id.* § 222.2. The code recognizes that robbery could be defined merely as a grade of theft but justifies its decision to define a separate offense “because of the special elements of danger commonly associated with forcible theft from the person.” *Id.* § 222.2 explan. note. By contrast, offense definitions that do not define one offense in terms of another do promulgate criminalization rules. Though many acts of forgery likely fall within the definition of theft by deception, *see id.* § 223.2, forgery does not require the commission of theft by deception, and thus forgery does make criminal certain actions that would not otherwise be crimes.

<sup>56</sup> 18 U.S.C. § 471.

<sup>57</sup> *Id.* § 503.

<sup>58</sup> *E.g., id.* § 471 (a fine or twenty years’ imprisonment); *id.* § 503 (a fine or five years’ imprisonment). Maximum authorized fines are set forth in a separate statutory provision. *See id.* § 3571.

<sup>59</sup> *See* 28 U.S.C. §§ 991–98.

The Guidelines provide a far more comprehensive grading scheme for federal crimes than the Model Penal Code does: based on certain enumerated characteristics of a particular violation of a particular statute, the Guidelines assign a quantitative point value between one and forty-three to each statutory violation, which is then adjusted based on certain factors that apply uniformly across all instances of a particular offense.<sup>60</sup> The resulting offense level determines, in part, what sentence the Guidelines authorize.<sup>61</sup> Thus, federal criminal law’s grading rules are split between statutory law and the Guidelines, with the latter employing a much more detailed grading scheme.<sup>62</sup> Statutes promulgate the federal criminalization rules that specify what conduct is a federal crime; the Guidelines promulgate the bulk of the federal grading rules that specify the seriousness of different kinds of criminal conduct.

As some of these examples indicate, a system of criminal law need not clearly divide its criminalization rules from its grading rules: a single legal provision may simultaneously promulgate both a criminalization rule and a grading rule. But regardless of how different bodies of law are structurally organized, criminalization rules and grading rules each correspond to a distinct function that must be performed somehow by all systems of criminal law, which must both define what conduct is a crime and must specify, for purposes of sentencing, the seriousness of each criminal act.<sup>63</sup> It is easiest to appreciate the difference between these two functions when the

---

<sup>60</sup> See U.S. SENTENCING GUIDELINES MANUAL §§ 1B1.1–.3 (U.S. SENTENCING COMM’N 2021) (explaining how to apply the Guidelines).

<sup>61</sup> See *id.* §5A. Originally, legislation required courts ordinarily to impose a sentence within the range the Guidelines authorized. See 18 U.S.C. § 3553(b)(1); *Mistretta v. United States*, 488 U.S. 361 (1989) (upholding the Guidelines’ constitutionality). Subsequently, the Supreme Court held that statutory provision unconstitutional. *United States v. Booker*, 543 U.S. 220 (2005).

<sup>62</sup> In fact, the grading rules statutes promulgate often bear very little relation to the grading rules promulgated by the Guidelines. Federal statutes distinguish different forgery offenses—and thereby grade forgery offenses—largely on the basis of the kind of writing forged. See *generally* 18 U.S.C. §§ 470–514. By contrast, the Guidelines largely disregard most such distinctions. Almost all forgery offenses fall under one of two guidelines. See U.S. SENTENCING GUIDELINES MANUAL app’x A, at 560. The kind of writing a defendant forges—in particular, whether the forged document is a bearer obligation of the United States—will determine which of those two guidelines applies. Compare *id.* § 2B1.1 (covering “Offenses Involving Counterfeit Bearer Obligations of the United States”) with *id.* § 2B1.1 (covering, among other offenses, “Forgery; Offenses Involving Altered or Counterfeit Instruments Other than Counterfeit Bearer Obligations of the United States”). Thus, the Guidelines do grade forgery to some extent according to the kind of writing forged. But this is the only such distinction they draw: neither section 2B1.1 nor section 2B5.1 further grades forgery according to the kind of writing forged. See *id.* §§ 2B1.1, 2B5.1 For the most part, then, the Guidelines simply disregard the grading distinctions made by statute, which affect only those Guidelines sentences that would otherwise exceed the statutory maximum. See *id.* §§ 5G1.1(a), 5G1.1 (c)(1).

<sup>63</sup> There is perhaps one exception: a system of criminal law that authorizes only a single possible sentence needs no grading rules specifying which criminal acts are more serious than others, since the sole available sentence will be imposed automatically on any offender judged guilty of an offense. (Although one could understand this



law itself distinguishes them by discharging each through different provisions of law, as the Model Penal Code and federal criminal law do (though in different ways). But even when the structure of legal provisions obscures the difference between how they define some kind of criminal conduct and how they specify the seriousness of different instances of that conduct, a logical distinction remains between the two: any body of criminal law must both criminalize conduct and grade crime, and it must contain criminalization rules and grading rules to do so, however it expresses them. That logical distinction, in turn, will be central to my analysis of the different ways in which criminal law may guide conduct.

## *2. Criminal Law and Deliberative Criteria*

In my analysis of individuals' decision-making, I argued that the rules governing their behavior may require them to employ either of two different kinds of criteria in deliberating over which action to perform. Each of those two kinds of criteria—decisive and non-decisive criteria—exerts a different kind of influence on deliberation. The distinction between the two, I will argue, corresponds to a difference between the criminalization rules and grading rules that are promulgated by criminal law. Each of these two sorts of rules aim to influence individuals' behavior, but they differ in how they do so: criminalization rules identify decisive criteria that individuals must employ in their deliberation, while grading rules identify non-decisive criteria. Consequently, differences in how decisive and non-decisive criteria exert their influence on deliberation may produce differences between the structure that criminalization rules and grading rules require to effectively perform their functions.

### *a. Criminalization Rules*

Decisive criteria influence behavior by specifying that certain otherwise available actions may not be chosen and therefore must be excluded from agents' deliberations. By contrast, non-decisive criteria influence behavior by providing

---

system to employ a grading rule that grades all crimes alike.) But so long as a system of criminal law authorizes the imposition of different kinds of sentences, it must somehow specify which crimes should receive which sentences.

agents with a reason for or against certain actions without overriding agents' own judgment as to whether that reason is outweighed by the balance of other applicable reasons. Criminalization rules exert the first of these two kinds of influence on individual behavior. When the criminal law defines a crime, its aim is not merely to indicate to individuals that such conduct is in some respect undesirable or that some reason exists for choosing alternative actions; criminal law does not ultimately defer to agents' own judgment about whether, all things considered, to violate the law. Rather, it specifies that particular actions are crimes in order to conclusively instruct individuals not to perform them. Indeed, the criminal law itself ordinarily makes clear the decisive nature of the guidance it provides. The Model Penal Code's purpose is "to forbid and prevent conduct";<sup>64</sup> the New York Penal Law's purpose is "[t]o proscribe conduct";<sup>65</sup> the California Penal Code defines a crime as "an act committed or omitted in violation of a law forbidding or commanding it."<sup>66</sup> These words— forbidding, proscribing, commanding, and preventing—all express conclusive judgments: to forbid, proscribe, command, or prevent an action is to decide whether or not others shall perform it, not to defer to their own judgments about whether or not to perform it. By defining conduct as a crime, the law conclusively settles whether individuals may perform it, by deciding that they may not. Those who are obedient to the law (which, to be sure, may not include everyone) cannot regard such actions as available options in deliberating about what to do. And excluding certain actions from deliberation is the characteristic function of a decisive criterion.<sup>67</sup> Thus, the criminalization rules that define when conduct is a crime purport on their face to promulgate decisive criteria for individuals to employ in decision-making.

Because decisive criteria govern a binary choice—namely, the choice of whether to treat a particular option as eligible or excluded in deliberation—a rule implementing a decisive criterion must divide actions into two groups, one containing actions to be excluded and the other containing actions to be treated as eligible. Criminalization rules possess this characteristic binary structure. Any criminalization

---

<sup>64</sup> MODEL PENAL CODE § 1.02(1)(a) (AM. LAW INST. 1962).

<sup>65</sup> N.Y. PENAL CODE § 1.05.

<sup>66</sup> CAL. PENAL CODE § 15. This definition follows Blackstone's. 4 WILLIAM BLACKSTONE, COMMENTARIES \*5.

<sup>67</sup> Raz explicitly argues that these mandatory norms impose second-order reasons. RAZ, *supra* note 40, at 58–59.

rule defines a particular sort of conduct as a crime—it specifies, say, that an individual commits forgery if, with a certain mental state, he alters or makes a writing in a certain way, or utters a writing altered or made in that way.<sup>68</sup> Thereby, the rule divides actions into two categories—namely, those that meet the definition set forth in the rule and those that do not. Offense definitions are supplemented by the provisions defining legal justifications, which qualify criminalization rules:<sup>69</sup> an action that falls within both an offense’s definition and a justification is grouped with actions that fall within neither, while actions that are criminalized but not justified form the other category on their own. This distinction between the two categories, in turn, corresponds to how the rule requires individuals to distinguish between those actions in their decision-making: actions that satisfy the offense definition but fall outside any justification must be excluded from deliberation, while all other actions are eligible to be chosen and may be included in deliberation. The law’s criminalization rules divide action into two categories—those that are criminalized, and those that are either not criminalized or both criminalized and justified—corresponding to an agent’s two possible alternatives concerning whether to admit a given action in deliberation. This binary structure further suggests that criminalization rules identify decisive criteria for agents to employ in deliberation.

Thus, taken on their own criminalization rules have the structure and objectives of rules that impose decisive criteria on deliberation. But criminalization rules do not merely to state rules of conduct; instead, they identify when individuals are liable to criminal punishment. And the functions of criminal punishment further support interpreting the rules identifying the limits of criminal liability to provide individuals with a decisive rather than non-decisive criterion to employ in deliberation. To be sure, criminal punishment has many functions, not all of which involve guiding individuals’ deliberations about which actions to perform.<sup>70</sup> But at

---

<sup>68</sup> MODEL PENAL CODE § 224.1(1) (AM LAW INST. 1962).

<sup>69</sup> *See generally* *Id.* §§ 3.01–10.

<sup>70</sup> In particular, on at least two standard views of the purpose of punishment, criminal punishment does not function by providing some criterion to guide individuals’ deliberation. First, one proposed purpose of punishment is incapacitation: imprisoning (or executing) dangerous offenders prevents them from committing the crimes that they might commit were they at liberty. *See, e.g.*, FRANKLIN E. ZIMRING & GORDON HAWKINS, *INCAPACITATION: PENAL CONFINEMENT AND THE RESTRAINT OF CRIME* 3 (1995). But, obviously, imprisoning individuals does not change their behavior by influencing how they choose from among the options available to them; rather, it changes their behavior by changing what options are available to them. Second, another proposed purpose of

least two such functions plausibly do guide individuals' deliberation: criminal punishment expresses a message of censure and condemnation of an offender's crime, and it deters other potential offenders by broadcasting the threat that lawbreaking will be punished. Each of these functions influences individuals' behavior by identifying a decisive criterion to be employed in deliberation.

According to expressivist accounts of punishment, the state punishes in order to express a negative attitude towards the action for which a defendant is punished. R.A. Duff describes punishment as "communicat[ing] the condemnation or censure that offenders deserve";<sup>71</sup> similarly, Gideon Yaffe argues that "by subjecting a defendant to harsh treatment . . . the state expresses its judgment of disapproval and condemnation of the defendant's act."<sup>72</sup> Advocates of this account take censure and condemnation to aim to persuade individuals to change their behavior: for Duff, for example, its goal is "that citizens recognize and accept the law's requirements as being justified and refrain from crime for that reason, or that offenders recognize the wrongfulness of their past crimes and refrain from future crimes for that reason."<sup>73</sup> Condemnation and censure, furthermore, exert a distinctive kind of influence. To condemn an action is not merely to criticize it in some respect without expressing any judgment about whether it should have been performed; rather, condemnation expresses the judgment that an action should not have been performed. As Duff says, condemnation aims to induce offenders to refrain from the behavior condemned, and refraining from acting reflects a conclusive judgment that the action should not be performed; condemnation does not merely identify one reason weighing against an

---

punishment is rehabilitation: to the extent that crime is understood to result from some sort of social, behavioral, or medical pathology that causes individuals to break the law, criminal punishment may produce law-abiding citizens by remedying that underlying pathology. *See, e.g.,* Nicola Lacey & Hanna Pickard, *From the Consulting Room to the Court Room? Taking the Clinical Model of Responsibility Without Blame into the Legal Realm*, 33 OXFORD J. LEGAL STUD. 1, 6–7 (2013). But the kind of correction envisioned by this approach generally aims to remedy the underlying cause of the defendant's criminal behavior, not to rationally guide the defendant's behavior. *See id.* at 6 ("At the sentencing stage, criminal acts and omissions tended to be regarded as symptoms of an underlying pathology which was assumed to contribute causally to the criminal misconduct. It was therefore natural to understand such behavior as by-passing the individual's cognitive and volitional capacities, belying the finding of criminality [sic] liability, and undermining the presumption of genuine individual choice or control: this is what it means for criminal misconduct to be understood as a manifestation of individual or social disease."). My focus in the text on the action-guiding functions of punishment, of course, is not meant to dismiss these other functions; they are simply not relevant to whether criminalization rules promulgate decisive or non-decisive criteria.

<sup>71</sup> R.A. DUFF, PUNISHMENT, COMMUNICATION, AND COMMUNITY 27 (2001).

<sup>72</sup> YAFFE, *supra* note 17, at 313.

<sup>73</sup> DUFF, *supra* note 71, at 81.

action while allowing that other reasons may nonetheless have been weightier. This sort of influence is characteristic of decisive criteria.

Deterrence theories of punishment instead take punishment to influence behavior through a mechanism other than persuasion. By subjecting those (and only those) who perform certain actions to undesirable treatment, the state deters individuals from performing those actions, since individuals may avoid that undesirable treatment by avoiding the actions that trigger its imposition.<sup>74</sup> On this account, deterrence rationally influences behavior by providing a consideration weighing against choosing certain options—namely, the existence of criminal punishment gives individuals a reason not to choose actions that make them liable to be criminally punished. By defining certain actions as crimes, criminalization rules identify which actions make an individual liable to be criminally punished. Thus, a criminalization rule provides the content of the criterion that deterrence introduces into individuals' deliberation: because individuals are liable to be punished only if they perform actions falling within the scope of a criminalization rule, individuals influenced by the threat of criminal punishment will respond in deliberation by disfavoring any action that falls within the scope of a criminalization rule.

What sort of criterion does deterrence provide—one that weighs decisively or non-decisively in deliberation? Criminal law is not the only body of law that attempts to influence individual behavior by imposing an undesirable consequence upon individuals who engage in some sort of socially harmful behavior—Pigouvian taxes, compensatory damages, and other legal remedies besides punishment may likewise alter individuals' choices by imposing an additional cost on those who select certain courses of action. All these forms of deterrence influence individuals' deliberations: the fact that individuals who perform certain actions suffer an undesirable consequence gives them a reason not to perform those actions. Thus, each kind of deterrence introduces a criterion into individuals' deliberation, one that favors actions that do not expose the actor to legal liability over actions that do and whose content thus depends on the principles defining the scope of liability. But there is no reason to

---

<sup>74</sup> Classical statements of this approach include CESARE BECCARIA, *ON CRIMES AND PUNISHMENTS AND OTHER WRITINGS* (Aaron Thomas ed., Aaron Thomas & Jeremy Parzen trans., 2008); and JEREMY BENTHAM, *AN INTRODUCTION TO THE PRINCIPLES OF MORALS AND LEGISLATION* (J.H. Burns & H.L.A. Hart eds., 1996).

think that these laws must all promulgate the same kind of criterion—either that all must promulgate decisive criteria or that all must promulgate non-decisive criteria. Why should deterrence always work the same way? Why would it either always aim to identify actions that individuals must exclude from deliberation entirely, or always aim merely to provide them with one reason not to choose that option that might nonetheless be outweighed by others depending on how the agent judges the balance of reasons? Rather, in my view laws may deter behavior by providing individuals with either a decisive or a non-decisive criterion disfavoring that behavior. Which kind of criterion the law provides depends on how it deters—in particular, on whether it fully or only partially deters individuals from performing the actions that give rise to liability.

Decisive and non-decisive criteria differ in how they influence an agent's deliberations. A non-decisive criterion has a weight of some particular strength, which determines the extent of its influence in deliberation. Agents who employ it are likelier to choose actions it favors and less likely to choose actions it opposes, but because its strength is limited agents will overrule it when sufficiently strong countervailing reasons exist. By contrast, weight is not relevant to the influence a decisive criterion exerts, for all such criteria demand the same response from individuals in their deliberation—namely, to exclude any option to which the criterion applies. Such a criterion is decisive because neither the strength of the reason it provides nor the strength of any potential countervailing reasons matter to how it influences decision-making. Instead, the only relevant question is whether the criterion applies to a particular option; settling that question settles whether or not it excludes that option from deliberation.

Thus, decisive and non-decisive criteria differ in how they function when they conflict with other reasons. Non-decisive criteria can be outweighed by countervailing reasons, in that sometimes an agent should overrule a non-decisive criterion and choose the option it disfavors, but a decisive criterion cannot be outweighed in this manner, since agents must always exclude options to which a decisive criterion (but no exception) applies. This difference in how the two kinds of criteria respond to conflicts, in turn, can serve as a test to identify which kind of

criterion a particular legal deterrent imposes, for the manner in which the law crafts a system of deterrence will determine the results of conflicts between those reasons and any countervailing reasons that may exist. The law deters by imposing undesirable consequences on those who perform certain actions, giving agents seeking to avoid those consequences a reason to avoid those actions. The strength of the reason the deterrent creates, in turn, depends on the severity of the consequence imposed: the more undesirable the consequence, the stronger agents' reason to avoid actions that might make them legally liable to suffer that consequence.<sup>75</sup> And the strength of the reason created by deterrence will influence how it behaves when it conflicts with other reasons that may apply. Thus, differences in the severity of the undesirable consequence a legal deterrent imposes may identify what kind of criterion it introduces into individuals' deliberation.

The strength of a non-decisive criterion, I have argued, is limited: though it exerts some influence on choice, in some contexts individuals rationally ought to overrule it and choose actions it disfavors, because stronger countervailing reasons exist. Thus, to impose a non-decisive criterion on deliberation a legal deterrent must create a reason of limited strength, which in turn requires the undesirable consequence it imposes to be sufficiently mild that individuals will choose to suffer it when sufficiently strong countervailing reasons exist. Indeed, the severity of the consequence must be precisely calibrated to give individuals a reason of the right strength to avoid actions that would make them liable: the strength of the reason created by deterrence must be exactly the same as the weight that individuals should grant the non-decisive reason when they employ it in deliberation. For, if the undesirable consequence is too mild, individuals will sometimes comply with countervailing reasons even when the non-decisive criterion should weigh more heavily, while if it is too harsh individuals will sometimes overrule countervailing reasons when they should outweigh the non-decisive criterion. Deterrence of this sort aims to deter only partially, imposing consequences mild enough that they only sometimes change the balance of an agent's reasons. And since the severity of the

---

<sup>75</sup> Of course, severity is only one consideration that influences the strength of the reasons created by deterrence.

consequences imposed determines how partially individuals will be deterred, that severity must be strictly limited lest it deter individuals too fully from acting.

Decisive criteria, however, are not of limited strength: if one disfavors an option (and no exception applies), agents must exclude that option from deliberation regardless of the existence of any countervailing reasons. Individuals ought to overrule a non-decisive criterion when sufficiently strong countervailing reasons exist; thus, laws imposing such criteria must deter sufficiently mildly that individuals will knowingly choose an action exposing them to legal liability if the reasons for doing so are strong enough. The deterrent must be mild enough it would be less undesirable to suffer it than to disregard adequately strong countervailing reasons, when they exist. But individuals should always exclude an option from deliberation when a decisive criterion applies. Consequently, a deterrent imposing a decisive criterion should not be intentionally designed to be weak enough that the reason it creates can sometimes be outweighed by countervailing reasons; rather, to impose a decisive criterion the law must deter harshly enough to create a reason disfavoring an option that outweighs whatever countervailing reasons apply in that case. If a deterrent imposes harsh enough consequences that an agent's reason to avoid them is stronger than his reasons to perform an action that would cause them to be imposed, then the fact that the deterrent applies will rationally settle whether he should perform the action. Consequently, he will have a decisive reason to avoid any actions that would make him liable to that consequence—just as he would were he to correctly apply a decisive criterion disfavoring those actions in his deliberation.

Such deterrence deters fully: it does not aim for individuals sometimes to disregard the reasons it creates, and thus it does not intentionally restrict the harshness of the consequences it imposes for the purpose of ensuring that individuals will sometimes act despite those consequences. Rather, it aims to impose a consequence harsh enough that it will always deter individuals from choosing actions that would impose them to liability. Of course, this does not imply that the consequences imposed must be maximally harsh; other considerations, ranging from mercy towards offenders to limited government budgets, may provide reasons to restrict the harshness of the sanctions imposed. Crucially, though, unlike with partial deterrence



this mildness is justified by objectives besides deterrence; the law does not limit its harshness so that it will fail to deter individuals whose countervailing reasons are sufficiently strong but rather because it is unwilling to bear the costs that would be required to deter fully. Like a system of partial deterrence, a system of full deterrence may confront contexts in which strong enough countervailing reasons do exist that individuals should perform actions that ordinarily would trigger the imposition of legal consequences. Systems of partial deterrence guide individuals to comply with those countervailing reasons in those contexts by ensuring that the strength of the deterrent is weak enough that it is outweighed by the countervailing reasons. But systems of full deterrence cannot take this approach, since they impose consequences designed to be severe enough to create reasons that cannot be outweighed. Instead, then, systems of full deterrence accommodate contexts such as these by creating exemptions to liability entirely: if no legal consequences are imposed whatsoever for acting when the countervailing reasons are sufficiently strong, then obviously nobody will be deterred from acting in those contexts.

Obviously, a legal deterrent can fully deter an agent's choices only if it is strong enough that an agent's countervailing reasons never outweigh it. This might appear to be a dissimilarity between full deterrence and decisive criteria: I have argued that strength is not a relevant parameter for decisive criteria, while clearly the strength of the reasons full deterrence creates are essential to its ability to deter fully. But although the importance to full deterrence of the strength of the reasons created by the consequence it imposes does clearly distinguish full deterrence from decisive criteria, there is nonetheless a clear similarity between the two: they each ensure that conflicts among reasons are resolved in the same way, though they may employ different mechanisms to produce those resolutions. If, that is, a legal deterrent is harsh enough that it gives individuals a reason to avoid liability that always outweighs any countervailing reasons, then the fact that a particular action would make the agent liable determines on its own what an agent should do, since the agent's reason to avoid liability outweighs any other reasons that might apply. Thus, though a full deterrent does create only an ordinary reason that must be weighed against all others to determine the balance of reasons, as is true of non-decisive criteria, it deters

severely enough that it simulates the effects on choice of a decisive criterion: in each case, that reason alone conclusively determines how an agent should choose.<sup>76</sup>

Which form of deterrence, then, does the criminal law employ? Certainly, the law does sometimes calibrate the severity of some consequences it imposes so that they deter only partially: a Pigouvian tax on emissions, for example, must be set low enough that regulated entities will emit, and pay the tax, when the benefits of doing so outweigh the harms of those emissions. But the criminal law does not take this approach in determining the severity of the punishments it imposes. Criminal punishments are not intentionally made milder so that the incentive they create not to offend is weaker than potential offenders' reasons to offend. Instead, criminal law takes the opposite approach: its punishments aim to be sufficiently harsh that they will outweigh whatever reasons potential offenders might have to offend. Bentham clearly articulates this principle in the first rule he proposes to govern the severity of criminal punishments: "The value of the punishment must not be less in any case than what is sufficient to outweigh that of the profit of the offense."<sup>77</sup> This principle, in turn, has guided actual systems of criminal law in determining the severity of the punishments they impose: criminal law aims to punish severely enough to fully deter potential offenders from crime.<sup>78</sup> Furthermore, criminal law handles decision-making contexts in which countervailing reasons do justify performing an action that otherwise would be criminal not by simply expecting individuals to disregard the deterrent, as individuals disregard a Pigouvian tax when their expected profits are substantial enough to pay the tax, but rather by defining exceptions to criminal

---

<sup>76</sup> Certain similarities exist between this analysis of how criminal punishment deters—namely, that punishment aims to induce offenders to have the kind of reason that would have led them to act correctly in their circumstances—and a recent analysis Gideon Yaffe has developed of deserved punishment for wrongdoing. GIDEON YAFFE, *THE AGE OF CULPABILITY: CHILDREN AND THE NATURE OF CRIMINAL RESPONSIBILITY* 103 (2018) ("[T]o give someone what they most deserve for past wrongdoing is to attach something to the act thanks to which the act has, for the agent, the very same reason-giving properties that it has, without that thing attached to it, to a better agent."). Yaffe does not connect this account of desert to the deterrent function of criminal punishment, however, as I do here.

<sup>77</sup> BENTHAM, *supra* note 74, at 166.

<sup>78</sup> United States Sentencing Commission, *Principles Governing the Redrafting of the Preliminary Guidelines*, 1(b) (1986); reprinted in Stephen Breyer, *The Federal Sentencing Guidelines and the Key Compromises Upon Which They Rest*, 17 *HOFSTRA L. REV.* 1, 47 (1988) ("The Guidelines seek to insure that all sentences convey the fact that crime does not and will not pay."); MODEL PENAL CODE § 1.02 cmt. 3(a) (AM. LAW INST. 1962) ("[T]he legislative purpose is that those who are convicted of defiance of the prohibition will be dealt with by a disposition that does not depreciate the gravity of the offense and thus imply a license to commit it.").

prohibitions that ensure no punishment will be imposed—namely, principles of justification.<sup>79</sup>

Criminal law, then, deters fully. If the punishment imposed for crime is severe enough to outweigh whatever reasons might weigh in favor of breaking the law, then the fact that an action falls within a criminalization rule on its own rationally determines that individuals should not perform it. To be sure, though the law aims to fully deter it does not always succeed at doing so: since individuals do sometimes violate the law, they must take themselves to have countervailing reasons that outweigh the disincentive created by the threat of punishment. But though criminal deterrence does not always succeed, it aims to give individuals a strong enough reason not to commit crimes that they may always exclude an action from deliberation just on the basis of its status as a crime. As I have argued, full deterrence corresponds to decisive criteria, in that full deterrence attempts to give individuals a reason that behaves the same way as decisive criteria do in resolving conflicts among reasons. Thus, considerations of deterrence, too, justify the conclusion the criminalization rules that define crimes promulgate decisive rather than non-decisive criteria for individuals to employed in deliberation.

*b. Grading Rules*

Once criminalization rules have defined which actions are crimes, grading rules specify the seriousness of different kinds of crime, thereby determining the sentences to be imposed upon different defendants. Because grading rules apply only to crimes, though, it may seem that they could not play any role in guiding action. After all, criminalization rules already instruct individuals how to deliberate about criminal actions: they are ineligible to be chosen and must be excluded from deliberation. And once an action has been excluded from deliberation, it is unclear what further guidance could remain for the law to provide. Once an individual has excluded crimes from her deliberations, as criminalization rules demand, criminal law leaves her free to choose among the options that are not crimes; it makes no further demands on her. What guidance, then, could grading rules provide? Of course, they

---

<sup>79</sup> See generally MODEL PENAL CODE §§ 3.03–.11.

could still perform an important function in determining the sentences imposed on offenders, but that function would merely influence how criminal cases are adjudicated, and would play no role in guiding conduct.<sup>80</sup>

Certainly, this objection is correct to some extent: there is no further guidance that the criminal law could provide to individuals who do comply with criminalization rules by excluding criminal acts from deliberation. But it does not follow that grading rules have no role to play in guiding individuals' conduct, for individuals do not always comply with the guidance criminalization rules provide concerning which actions are eligible to be performed. Instead, individuals sometimes decide to engage in crime. Once an option is excluded from deliberation, there is no further guidance for the law to provide, but the law may further guide individuals who have chosen to commit a crime: they will include some criminal acts in their deliberation, and the law could provide them with guidance concerning which criminal acts to perform. Furthermore, the law often should provide that guidance. In Beccaria's words, "it is in the common interest not only that crimes not be committed, but that they be rarer in proportion to the harm that they do to society."<sup>81</sup> Thus, the law often should guide offenders to choose offenses that are likely to cause less harm to the social interests protected by criminal law.

It will be rare that on a single occasion the options available to an individual will include every kind of crime, of course: it is difficult to imagine a context in which an individual might deliberate, say, between either importing a zebra mussel into the United States<sup>82</sup> or violating the privacy of a child witness.<sup>83</sup> But there are many contexts in which individuals do choose among different criminal actions—

---

<sup>80</sup> In his taxonomy of the criminal law, for example, Paul Robinson treats the law's grading function as wholly distinct from its rule-articulation function, the latter of which is responsible for guiding individuals' behavior. PAUL ROBINSON, *STRUCTURE AND FUNCTION IN CRIMINAL LAW* 125 (1997); Paul H. Robinson, *A Functional Analysis of Criminal Law*, 88 NW. U. L. REV. 857, 857 (1994). Rather than guiding the conduct of ordinary citizens, on this view, grading rules guide only the decisions of judges adjudicating cases. For analysis of this distinction in the audience to which different rules of criminal law are directed, see generally Meir Dan-Cohen, *Decision Rules and Conduct Rules: On Acoustic Separation in Criminal Law*, 97 HARV. L. REV. 625 (1984); and Paul H. Robinson, *Rules of Conduct and Principles of Adjudication*, 57 U. CHI. L. REV. 729 (1990).

<sup>81</sup> BECCARIA, *supra* note 74, at 165; see also BENTHAM, *supra* note 74, at 165 ("But if a man must needs commit an offence of some kind or other, the next object is to induce him to commit an offence less mischievous, rather than one more mischievous: in other words, to choose always the least mischievous, of two offences that will either of them suit his purpose.").

<sup>82</sup> See 18 U.S.C. § 42(a)(1) (2018).

<sup>83</sup> See *id.* § 403.

whether to forge an obligation of the United States or some other writing,<sup>84</sup> whether to steal goods from their owner's person by threat or in some other way,<sup>85</sup> or whether to assault another with or without a deadly weapon.<sup>86</sup> Some of these offenses are clearly more threatening to the interests of society than others. Furthermore, the fact that an individual has disregarded a criminalization rule need not imply that he will disregard grading rules, too. If multiple different crimes may all achieve the objective motivating an individual to break the law, he might follow its guidance in choosing among those crimes despite disregarding its guidance in being willing to perform them at all. And since by definition unintended results are not part of an agent's reasons for acting, offenders may choose actions less likely to produce unintended harms without undermining their reasons for offending. Thus, a role does remain for grading rules to play in guiding conduct: they may direct individuals choosing between multiple criminal acts to choose crimes that are less harmful.

Obviously, in some sense this kind of guidance represents a second-best outcome: individuals ideally would exclude actions defined as crimes from their deliberations altogether, and guidance as to which crimes are relatively less harmful is relevant only if individuals will not choose the optimal course of action. Thus, there may appear to be some internal tensions in the guidance the law communicates: it both instructs individuals not to perform certain actions and tells them how, should they ignore that instruction, they should choose among the actions they should not perform.<sup>87</sup> But although there may seem something odd about instructions of this sort, there is nothing incoherent or irrational in this guidance. The law is genuinely committed to both of its components: society's interests are best advanced if individuals exclude criminal acts from their deliberations entirely, but failing that society's interests are better advanced if individuals commit less harmful crimes.

In some cases, to be fair, one might worry that combining these two instructions undermines the guidance they provide—that, in particular, rules stating

---

<sup>84</sup> Compare U.S. SENTENCING GUIDELINES MANUAL § 2B1.1 (U.S. SENTENCING COMM'N 2021) with *id.* § 2B5.1.

<sup>85</sup> See MODEL PENAL CODE § 223.1(2)(b) (AM. LAW INST. 1962).

<sup>86</sup> See *id.* § 211.1(2)(b).

<sup>87</sup> Matthew Mandelkern has analyzed the intuitive strangeness of the related phenomenon of giving an order while acknowledging that it might not be followed, as, for example, in the sentence "You must close your door, but I don't know whether you will." Matthew Mandelkern, *Practical Moore Sentences*, 55 *NOÛS* 39 (2021).

which crimes are preferable to others somehow weakens the strength of rules prohibiting crime entirely.<sup>88</sup> But the manner in which grading rules operate limits the risk that they will undermine criminalization rules. They do not exempt less serious offenses from punishment, which surely would undermine the force of a rule criminalizing those offenses by seeming to communicate, inconsistently, both that individuals must exclude those options from deliberation and that those options nonetheless may be chosen. Instead, grading rules increase the severity of sentences imposed for more serious offenses. All crimes are punished; more serious crimes are punished more severely. This guidance, in effect, is that individuals must not perform any crimes, especially not certain more serious ones. There is nothing inconsistent or self-defeating about it.

This account of how grading rules influence individuals' conduct makes clear which kind of criterion they promulgate. Grading rules guide conduct by influencing individuals' deliberations about which criminal act to perform. They exert influence, then, only after an individual has already disregarded the law's criminalization rules by deciding to perform a criminal act of some sort. By disregarding the law's criminalization rules, in turn, such individuals disregard the decisive criteria those criminalization rules promulgate: while the law instructs individuals to exclude criminal acts from their deliberation entirely, such individuals instead treat such acts as eligible to be performed and deliberate about which to choose. Grading rules, then, influence the conduct only of those individuals who disregard the law's decisive criteria. For them to promulgate a further decisive criterion would be pointless: grading rules guide only individuals who will not heed that criterion. Instead, if grading rules are to influence conduct at all, they must do so by promulgating a different kind of criterion—namely, a non-decisive one. Decisive and non-decisive criteria each influence different stages in the deliberative process: the former guide individuals in including certain actions as available options in their deliberations, then the latter subsequently guide individuals in evaluating the options between which they must deliberate and choose. Because grading rules are relevant only to

---

<sup>88</sup> Mandelkern gives an explanation essentially along these lines to analyze the strangeness of the sentences he considers: in his view, to give an order involves placing oneself in a position of authority, which is undermined by the acknowledgement that one's own order might not be followed. *Id.* at 54–55.

individuals who have already disregarded the law's guidance at the first stage of this process by including criminal acts as available options in deliberation, only the second stage remains in which grading rules may influence behavior. Thus, if they are to influence behavior at all, they must promulgate the sort of criterion that guides individuals in comparing multiple options to determine which the balance of reasons supports—namely, a non-decisive criterion.

Furthermore, it is precisely because grading rules promulgate a non-decisive rather than a decisive criterion that there is no internal contradiction between the law's criminalization rules and grading rules. If grading rules promulgated a decisive criterion, it would obviously have to be a different decisive criterion than is promulgated by criminalization rules: since grading rules draw distinctions only among actions that are crimes in order to guide individuals in deliberating among different criminal acts, they would instruct individuals to exclude only some criminal acts from deliberation. It would be pointless for grading rules merely to recapitulate criminalization rules that individuals have already decided to disregard. And if criminalization rules and grading rules promulgated distinct criteria, there would exist actions that are defined as crimes and thus ruled out by criminalization rules but that are not ruled out by grading rules. This combination would produce genuinely contradictory instructions: criminalization rules would tell individuals to exclude certain actions from their deliberations, while grading rules would then undermine that instruction by telling individuals that they may in fact include those actions. This contradiction disappears, by contrast, once grading rules are understood as promulgating non-decisive criteria. Grading rules do not undermine criminalization rules' instruction to exclude certain options because they do not purport to instruct individuals as to which options to include in deliberation at all; instead, they instruct individuals how to compare the actions between which they are deliberating. Because they address an entirely different question than criminalization rules, grading rules do not undermine the answers that criminalization rules give.

Finally, grading rules possess the structure associated with non-decisive rather than decisive criteria. In general, grading rules operate by introducing a ranking scheme that defines different degrees of wrongdoing, identifies the degree to which

each criminal act is wrongful, then specifies what sentence should be imposed for offenses of each degree. The Model Penal Code identifies each criminal act as a petty misdemeanor, misdemeanor, or felony of the first, second or third degree,<sup>89</sup> then specifies what sentences should be imposed on offenders whose crimes fall into each category;<sup>90</sup> the Federal Sentencing Guidelines define the offense level of each act violating federal criminal law,<sup>91</sup> then specify what sentence should be imposed on an offender given the level of his offense.<sup>92</sup> (Other factors also influence the ultimate sentence, of course.) These rankings compare offenses: offenses are worse than others to the extent that they are punished more severely. Whether these comparisons are understood as expressing greater censure of<sup>93</sup> or as providing a stronger deterrent for<sup>94</sup> those offenses that are punished more severely, they clearly identify some offenses as worse than others. Non-decisive criteria influence deliberation by giving individuals some basis with which compare their options in deliberation; by identifying some offenses as worse than others, then, grading rules identify a non-decisive criterion that guides individuals choosing between offenses to select those that are punished less severely.

Decisive criteria guide conduct by classifying offenses into two categories; individuals must not choose options excluded by the criterion, regardless of any other applicable reasons in their favor. But grading rules make no such binary classification. The sentencing guidelines, for example, identify forty-three different offense levels; they do not further specify any binary division among them.<sup>95</sup> Thus, there is no category of actions that grading rules instruct individuals to exclude from deliberation. Furthermore, assigning a particular offense level to a particular criminal

---

<sup>89</sup> *E.g.*, MODEL PENAL CODE § 233.1(2) (defining grades of theft).

<sup>90</sup> *Id.* §§ 6.03, .06, .08.

<sup>91</sup> *E.g.*, U.S. SENTENCING GUIDELINES MANUAL § 2B1.1 (U.S. SENTENCING COMM'N 2021) (defining the offense levels of different property offenses).

<sup>92</sup> *Id.* § 5A (providing the sentencing table for sentences of imprisonment).

<sup>93</sup> DUFF, *supra* note 71, at 132 (“We must determine not just that an offender deserves censure but how severe that censure should be: the more serious the crime, the more severe the deserved censure. . . . Thus the severity of the penal hard treatment will communicate the severity of the censure: the more severe the hard treatment, the more severe the censure it communicates.”).

<sup>94</sup> BECCARIA, *supra* note 74, at 17 (“Therefore, the obstacles that deter men from committing crimes must be more formidable the more those crimes are contrary to the public good and the greater are the incentives to commit them.”); BENTHAM, *supra* note 74, at 168 (“Where two offences come in competition, the punishment for the greater offence must be sufficient to induce a man to prefer the less.”).

<sup>95</sup> *See* U.S. SENTENCING GUIDELINES MANUAL part 5A (U.S. SENTENCING COMM'N 2021).



act does not, on its own, determine how individuals should treat that act in their deliberations. The fact that a particular crime's offense level is fifteen, for example, would be a reason to choose it if one's other options are all of higher offense levels, but would be a reason not to choose it if one's other options are all of lower offense levels. While decisive criteria in themselves determine whether an act is eligible to be chosen, then, grading rules do not; instead, whether an individual should choose to perform an offense graded at a particular level depends on whether the balance of reasons favors that choice, given the other options available. Consequently, grading rules cannot plausibly be interpreted as promulgating a decisive criterion. Criminalization rules promulgate decisive criteria that exclude all actions defined as crimes; grading rules subsequently promulgate a non-decisive criterion that guides individuals intending to offend in choosing among different offenses.

### *C. Guiding Action by Its Results*

Individual deliberation, I have argued, employs two kinds of criteria, decisive and non-decisive, that each influence choice in different ways. The criminal law, in turn, contains two kinds of rules that perform two distinct functions, criminalization and grading; the two types of rule, respectively, promulgate decisive and non-decisive criteria. This distinction between how criminalization rules and grading rules influence conduct, I will argue, explains why results may legitimately play different roles in criminalization than in grading. Decisive criteria influence conduct by classifying an agent's options in two categories and instructing agents to exclude from deliberation all actions that fall into one, and criminalization rules cannot effectively identify such a criterion by defining crimes in terms of results alone. By contrast, non-decisive criteria merely compare the merit of different actions without identifying any category of actions that must always be excluded from deliberation; grading rules can identify such a criterion by grading offenses by results alone. Thus, I will conclude, the law must require mens rea as a minimum condition of criminal liability, but it may grade offenses by their results regardless of whether the defendant possesses some type of mens rea with respect to those results.

### *1. Criminalization by Results*

By defining offenses, criminalization rules divide actions into two categories—those that are and are not crimes. Thereby, I have argued, they promulgate a decisive criterion that identifies which actions individuals must exclude from deliberation. Broadly speaking, two kinds of relationship might exist between the criminalization rule and the decisive criterion. First, a criminalization rule might directly identify a decisive criterion: a criminalization rule might instruct individuals to exclude an action from deliberation just in case that action is a crime. Second, a criminalization rule might somehow indirectly identify a decisive criterion: that is, by identifying one set of actions as crimes, it might instruct individuals to exclude a different set of actions from their deliberation. A criminalization rule that defines an offense in terms of results alone cannot effectively define a decisive criterion using either of these two approaches. Consequently, criminalization rules cannot define offenses in terms of results alone. Instead, in order to promulgate a decisive criterion, criminalization rules must define offenses in terms of the mental states of the defendant—precisely as they do when they employ *mens rea*.

#### *a. Excluding Harms*

When deliberating about what to do, individuals form some sort of mental representation of the various options between which they must choose. These representations each rely on the beliefs individuals possess about what would happen were they to pursue each of the actions between which they are deliberating. Someone deliberating, say, about whether to run to catch a subway might have beliefs about the train's location, whether it is still discharging passengers, how soon it is likely to leave, how quickly she could reach the nearest door, and whether she will ultimately make the train. These beliefs do not represent the world perfectly, however. We are not omniscient; sometimes we are not sure about some fact and must suspend judgment, and other times our beliefs are simply mistaken. In evaluating their options during deliberation, though, individuals can employ only the information that they actually possess. Any criterion, decisive or non-decisive, favors actions possessing some particular property over actions that lack it, but individuals

can apply that criterion to their options only by favoring actions that they believe possess the property over actions that they believe do not. When an individual's beliefs are mistaken or incomplete—if she does not recognize that an action possesses a property it does in fact have—she nonetheless must apply the criterion based on those imperfect beliefs rather than based on the correct facts about the world, for she lacks access to those facts except through the beliefs she forms about them. An individual might, for example, accept a criterion that favors running for the train only if by running she would catch it. But in applying that criterion to make any particular decision about whether to run for the train, her decision will depend not on whether by running she would catch it in fact but rather on her beliefs, such as her belief about whether by running she would catch it. The world may actually be different from how it seems from our point of view, but we must nonetheless make decisions from that point of view.

Because our representations are imperfect, the actual state of the world may vary in ways that are not captured by our beliefs about it. Sometimes an individual running to catch the train will successfully catch it; other times she will not. Of course, there will frequently be differences between the beliefs that individual possesses in the former kind of situation and those she possesses in the latter: for example, most people probably believe they will catch the train by running more frequently on occasions when they actually will than on occasions when they actually will not. But sometimes differences in facts about the world will not be reflected by differences in our beliefs about it: on different occasions, two individuals may share the same set of beliefs even though different facts are true on those different occasions. Sometimes an individual running to catch the train actually will catch it; sometimes an individual with the same beliefs about her circumstances will fail to actually catch it. These possibilities will only rarely be equally likely: someone quite close to a train still discharging passengers will probably make it, while someone at the other end of the platform probably will not. But each outcome is possible. An individual's beliefs about what would result from choosing a particular action do not usually themselves make a particular fact about those results necessary or impossible; instead, any particular set of beliefs is generally compatible with the possibility that

any particular fact about the world will be true or false (including, obviously, beliefs about whether that fact is true or false). Distinct possibilities exist in which individuals' beliefs about some available course of action are identical in each possibility, yet choosing it would produce different results in some such possibilities than in others.

Because in deliberating about our options we can employ only the information we actually possess, then, the existence of distinct possibilities in the world that are each compatible with the same set of beliefs restricts the kinds of distinctions we can draw among our options in deliberation. A decisive criterion, I have argued, influences deliberation by classifying each option into one of two categories: for any given option, it specifies whether or not that option is to be excluded from deliberation. But if we deliberate based only on our beliefs about our options, then a decisive criterion can classify an option as excluded or not based only on the beliefs that we have about it. Because a decisive criterion must determine whether or not to exclude an option based on the agent's beliefs alone, in turn, its guidance must apply in the same way whenever she has those beliefs. Either she must always exclude an option given those beliefs, or not: the criterion cannot require her sometimes to exclude an option and permit her sometimes to include the very same option about which her beliefs are identical, for absent any differences in what properties she believes that action to have, the criterion has no way to indicate when exclusion is required and when it is not.<sup>96</sup>

---

<sup>96</sup> Strictly speaking, a third option exists: she might exclude or include the option on different occasions according to some fixed percentage—she could, say, exclude it only one-third of the time—then use a random process, with appropriate odds, to decide whether to exclude or include on any particular occasion. Options of this sort are quite important in game theory, which terms them mixed strategies. Don Ross, *Game Theory*, STAN. ENCYCLOPEDIA PHIL. 3 (Edward N. Zalta ed., Fall 2021 ed.), <https://plato.stanford.edu/archives/win2019/entries/game-theory/>. But this qualification does not undermine my argument. A decisive criterion must apply equally to all options characterized by the same beliefs even when mixed strategies are allowed: they simply must require the same mixed strategy whenever the agent faces an option characterized by those beliefs. And, of course, though a mixed strategy does entail that an agent will sometimes act differently on different occasions when her beliefs are the same, she does not choose different actions on those occasions based on facts not reflected in her beliefs; rather, she chooses based on some randomization process whose results, presumably, must be reflected in her beliefs. (Indeed, one might think that the way a randomization process works is by adding beliefs that distinguish the occasions from one another—if, say, the agent decides by rolling a die and excluding the option when she rolls a one or a two, then it is no longer true that her beliefs are identical on occasions when she excludes it and when she does not, for she will have different beliefs about the results of her roll of the die.) Furthermore, even if mixed strategies are an option in theory, it does not seem that the criminal law ever requires individuals to act by employing a mixed strategy.

In particular, if some set of beliefs about an option is compatible with a particular fact about the world being either true or false, the decisive criterion must nonetheless classify that option in the same manner regardless of whether the fact actually is true or false, for its actual truth or falsity is not reflected in her beliefs and thus cannot influence her deliberation. Suppose, for example, that someone is deliberating whether to run for the train given some particular set of beliefs about its location, her speed, how soon it will depart, whether she will make it, and so forth. Out of all the occasions in which she has that exact set of beliefs, sometimes she will make the train by running and sometimes she will not. But because that difference is not reflected by some difference in her beliefs, which are the same on all such occasions, any decisive criterion must classify all such options alike in her deliberation: it must either instruct her always to exclude running from deliberation given those beliefs, or not. It cannot instruct her whether or not to exclude running based on differences in those actions about which she lacks information.<sup>97</sup>

If the criteria we employ in deliberation cannot distinguish between actions based on information we lack, though, individuals cannot employ a decisive criterion requiring them to exclude an option based not on their beliefs or other mental attitudes about it but rather based only on whether some state of affairs will in fact obtain if they choose it. For example, individuals cannot exclude an action from their deliberation if and only if that action will in fact cause the death of another. Individuals exclude options in deliberation based on their beliefs about those options, but two actions believed by agents to be identical will sometimes differ in whether they actually cause death. Because that difference is not reflected in the agent's beliefs, agents must treat the two options identically in deliberation: they cannot exclude only some based on a difference not captured by their beliefs, since—precisely because that difference is not captured by their beliefs—they would have no

---

<sup>97</sup> To be sure, the information an individual possesses need not be fixed: individuals may sometimes acquire additional information by inquiring before they act, and with that additional information they may become able to distinguish among different options that appeared identical prior to inquiry. In chapter 3 of this dissertation, I analyze criminal negligence as a violation of the rules governing when individuals must inquire before acting. But, obviously, inquiry at some point must end, and when it does our information will still be imperfect. And at that point—when an individual must decide what to do—she cannot employ a criterion that requires her to make distinctions based on information she does not possess. (And the rules that govern when she must inquire further cannot themselves identify when inquiry is required in terms of information that she lacks.)

way of identifying which ones to exclude. In order to employ a strategy in deliberation that would require excluding a particular action that would in fact cause death, individuals must adopt a criterion that excludes any action about which they possess identical beliefs, even though some such actions may not cause death. Conversely, in order to employ a strategy in deliberation that would not exclude a particular action that would not, in fact, cause death, individuals must adopt a criterion that does not exclude any action about which they have identical beliefs, even though some such actions may cause death. If there is no difference between an agent's beliefs on two occasions, the agent has no basis upon which to treat an action differently in deliberation on those occasions, even if the results of performing that action would in fact be different. And if individuals cannot in fact employ a decisive criterion that excludes actions based on their actual results, a criminalization rule that criminalizes actions based on their actual results could not impose a decisive criterion that directly excludes actions that violate the rule, since individuals could not actually employ it.

*b. Excluding Risks*

A criminalization rule framed in terms of results alone, then, cannot influence individuals to exclude options from deliberation just in case they would constitute crimes, since individuals cannot deliberate using a criterion that distinguishes actions based on information that they do not possess. Instead, any decisive criterion must exclude options from deliberation based on what an individual believes about them, and any criminalization rule must therefore guide individuals to exclude options from their deliberation just in case they believe those options to have some particular feature. For a criminalization rule framed in terms of results alone to guide conduct effectively, then, that rule must somehow specify what beliefs about a particular option would require an agent to exclude that option from deliberation, thereby identifying a decisive criterion that excludes options based on the agent's beliefs about them. That criterion could not exclude options just in case they constitute crimes, at least given a criminalization rule on which whether an action constitutes a crime depends only on its actual results, not on the agent's beliefs about it. But

perhaps a criminalization rule defined in terms of results alone could indirectly identify the options individuals must exclude from deliberation.

Because of the manner in which criminal punishment influences behavior, however, I will argue that criminalization by results cannot indirectly identify a decisive criterion. A decisive criterion classifies an individual's options into two categories in order to identify which must be excluded from deliberation and which may instead be included. For a criminalization rule to identify a decisive criterion indirectly, then, it must somehow communicate to individuals which options may be chosen in deliberation. In general, though, because human knowledge is imperfect there always exists some possibility that choosing any option available in deliberation will cause a particular result, even if it is quite unlikely to do so. Consequently, if a criminalization rule defines an offense in terms of results alone, thereby making an individual's criminal liability depend only on the results of his conduct, individuals will always face some risk of criminal liability, no matter which option they choose, since there will always be some possibility that by choosing any particular action they will cause the prohibited result and thereby commit a crime. But precisely because criminal punishment provides individuals with a decisive criterion regardless of whether its aims are interpreted as expression or deterrence, even the risk of criminal liability suffices to exclude an option from deliberation. Criminalizing conduct by its results, then, does not identify a distinction between those options that must be excluded from deliberation and those that may be included; instead, it effectively instructs individuals to exclude all options from deliberation. And since no agent can follow that guidance—some option must be chosen, even if it is simply the option to do nothing—criminalization by results thereby fails to effectively guide individuals' conduct.

On expressivist accounts, punishing individuals who perform certain actions communicates the state's condemnation of those actions. This condemnation, in turn, guides conduct decisively: the law's condemnation of a particular action communicates that individuals must exclude such actions from their deliberation. A criminalization rule that defines an offense based on the results of conduct alone, without requiring *mens rea*, would punish and condemn any individual whose

conduct actually does cause a particular result. In so doing, it would condemn the action he performed as wrongful, thereby communicating that he should not have performed it. As I have argued, though, in deliberation individuals cannot distinguish between different actions based on differences between them that are not reflected in differences between the individual's beliefs about them: given identical beliefs about two possible actions that differ in some respect, individuals cannot treat those actions differently in deliberation. Consequently, if individuals must exclude from their deliberations a particular action that in fact causes a particular result—a requirement the law communicates by punishing and condemning individuals who do in fact choose it—they must likewise exclude all options about which they have identical beliefs, for given those identical beliefs there is no basis upon which they might distinguish between those options in order to exclude only some.

Consider, then, an individual deciding whether to include a particular action as an option in deliberation, given a criminalization rule that criminalizes conduct based on its results alone. Because of the possibility of error in his beliefs, there is always some possibility, however low, that choosing this option will cause a result in virtue of which it will constitute a crime. By punishing individuals when they do actually cause that result, the law conveys its condemnation of their action. This condemnation, furthermore, does not merely indicate that the action performed is in some respect worse than others; rather, it conveys that the action was wrong and should not have been performed. Thus, by punishing the law makes a decisive determination about the option the agent chose: it is not eligible to be chosen and, therefore, the agent must exclude it from deliberation. But individuals have no way to distinguish that option, which did in fact cause a prohibited result, from any other option about which they possess identical beliefs. Thus, they cannot exclude the option only when it will, in fact, cause that result; rather, they can exclude it in those cases only by always excluding it.

This argument, though, applies to every option that an individual might face. Because the possibility of error in one's beliefs always exists, though it may be quite unlikely, there is always some possibility that any particular action will cause a prohibited result. If the law criminalizes actions simply based on whether they cause



that result, then no matter what an individual believes about an action, it will be punished if it causes a particular result. That punishment communicates that individuals must exclude the action from deliberation in that case. But they can exclude it from deliberation in that case only by adopting a criterion that excludes every action about which they have identical beliefs from deliberation. Since this argument would apply to any action, a criminalization rule framed in terms of results would require individuals to exclude all actions from deliberation. But individuals must choose some option available to them; they cannot exclude all options from their deliberation. This guidance, then, is no guidance at all: a criminalization rule framed in terms of results alone cannot effectively guide conduct.

In addition to expressing condemnation, criminal punishment deters individuals from criminal activity: the threat of being punished for an action gives individuals a reason to avoid it. Just as criminalization rules cannot effectively employ results alone because of the decisiveness of the condemnation criminal punishment expresses, though, such rules similarly cannot deter effectively because criminal punishment—unlike other legal deterrents—aims to deter conduct fully rather than only partially. In particular, criminal punishment aims to be harsh enough to give individuals a reason to avoid committing crimes that is strong enough to outweigh whatever countervailing reasons they may have in favor of criminal conduct. Thus, a rule criminalizing conduct based on its results alone must punish severely enough to outweigh an individual's reasons for performing the action that produced that result. If an individual's beliefs about two actions are identical, though, the reasons he recognizes for performing the two actions must be the same: any reason to perform either action must depend upon some fact about that action, and since there is no difference in the facts about each action of which he is aware there can be no difference in the reasons of which he is aware, either. Thus, if a criminalization rule deterred individuals from performing actions that actually would cause a prohibited result, it would also deter individuals from performing any other action about which they had identical beliefs. And since it is always possible that choosing any option might actually cause any result, given the possibility that our beliefs might be mistaken, this deterrence would therefore encompass every option

individuals face whatsoever. If criminal punishment were to deter fully, criminalization by results would deter individuals from choosing any of the actions available to them.

Of course, no criminalization rule could actually succeed at deterring in this way. In practice, individuals must include some options in deliberation: nobody can simply refuse to act but rather must, at some point, perform some action. A system of criminal punishments designed to fully deter individuals from choosing to perform any action at all cannot succeed; instead, the law must accept that individuals will include some options in deliberation. But if criminal law cannot exclude all options, it would be counterproductive to define criminalization rules that aim to give individuals decisive reason to exclude all options from deliberation. Criminal law is extremely costly: criminal punishment gives individuals reasons to modify their conduct by imposing severe punishments that substantially harm the individuals punished. The law should not incur these high costs for nothing; it should not impose punishments that cannot achieve their aims. Because individuals cannot exclude all options from deliberation—at some point they must choose to do something—at least one option must exist that cannot be excluded from deliberation by the threat of criminal punishment. Therefore, punishing individuals who choose that option would be pointless, incurring the high costs of punishment even though deterrence in this case is impossible, since some option or other must in any event be included in deliberation. In fact, the law can guide conduct more effectively by exempting individuals who choose sufficiently good options from punishment: if choosing better options allows individuals to avoid the risk of being punished should they happen to cause a particular result, then they will be more likely to choose a better option.

Whether one focuses on its expressive or its deterrent function, criminal punishment aims to influence conduct decisively. Criminalization by results runs aground precisely because of the decisiveness of the reasons criminalization rules provide. Because we employ our beliefs in deliberation, we cannot directly govern our choices according to the actual results of our actions; rather, we deliberate based on our beliefs about the likely results of choosing each option, or about the aspects of each option that make certain results more or less likely. And because the likely

results of different options differ, the value of a particular result provides reasons of different strengths that apply to different actions depending on how likely they are to produce it. In particular, the strength of the reason is its value, discounted by the likelihood that the result will not occur, in which case acting will not incur the disvalue associated with that result. Because different actions differ in their likelihood of causing particular results, those results must give rise to reasons that weigh in deliberation with different strength. They must provide non-decisive rather than decisive criteria. And because they provide non-decisive criteria, such a reason can apply to every option available to a particular agent: non-decisive criteria can be outweighed, and thus their applying to all options does not entail that no option may be chosen.

By contrast, criminal punishment gives individuals decisive reason to avoid criminal conduct. The rational force of a decisive reason does not depend on its strength: individuals must treat any such reason as ruling out the actions to which it applies, regardless of whatever countervailing reasons might exist. If a decisive reason applies to multiple options in deliberation, then, it must exert the same kind of influence on each option, since there is only a single kind of influence that such a reason can exert. Thus, while reasons of a particular weight can be discounted—say, by the likelihood of a result's not occurring—decisive reasons cannot. If they apply to an option at all they apply by excluding it from deliberation; the rational force of a decisive reason cannot be reduced while the reason still applies but can only be removed entirely by fully exempting an option from the reason's scope. Thus, while reasons grounded in the results of acting can apply to all an individual's options, weighted differently, decisive reasons must not apply at all to at least some options if they are to guide conduct, for a rule that excludes all options leaves individuals with no basis upon which to decide what to do. Because of this mismatch between the binary structure of decisive reasons and the graded structure of probabilities, criminalization rules that employ results alone cannot guide conduct effectively.

Instead, criminalization rules must exempt some options from the scope of the decisive criteria they promulgate: rather than identifying only a forbidden result, which would weigh to some extent against choosing any option, that criterion must

divide the options to be excluded from deliberation from the options an individual is permitted to choose. To mark this division, the criminal law requires mens rea for guilt. Criminally punishing only those defendants who possess a particular mental state—a particular belief, say, or the awareness of a risk—identifies the options individuals must exclude in terms of the mental attitudes they possess towards those options.<sup>98</sup> So long as an individual does not possess the requisite mental state—does not believe or perceive a risk that by choosing some option he will perform the actus reus of a crime—he is not liable to criminal punishment for choosing that option. The criminalization rule thereby marks a binary division among an individual's options, classifying them into two categories. That classification identifies the decisive criterion individuals must employ in deliberation: individuals must exclude options towards which they possess the mental state identified in the definition of the relevant grade of mens rea; all other options may be included.

On my account, criminalization rules guide conduct effectively only because they incorporate some kind of mens rea into their definitions of criminal offenses. By contrast, many influential accounts of the structure of criminal law deny that the law's incorporation of mens rea plays any role in guiding individuals' behavior.<sup>99</sup> Paul Robinson, for example, takes the articulation of rules of conduct and the definition of the minimum conditions of liability to be distinct functions of the criminal law;<sup>100</sup>

---

<sup>98</sup> Because criminal negligence is defined in terms of the risks that a reasonable individual would perceive in the defendant's circumstances, not in terms of the defendant's actual mental states, criminal negligence might be understood as a form of mens rea on which a defendant's guilt does not depend on his mental states. On such interpretations, criminal negligence would be in tension with the account I defend here. In chapter 3 of this dissertation, however, I defend an account of criminal negligence according to which whether a defendant is criminally negligent does depend on whether he possesses a particular mental state.

<sup>99</sup> To be sure, some accounts of mens rea do recognize its function in guiding behavior. In particular, a number of scholars, most notably H.L.A. Hart, have emphasized that mens rea gives individuals notice of which actions the law deems criminal and, consequently, will punish. See H.L.A. Hart, *Punishment and the Elimination of Responsibility*, in *PUNISHMENT AND RESPONSIBILITY: ESSAYS IN THE PHILOSOPHY OF LAW* 158, 181 (2d ed. 2008) (“[A] primary vindication of the principle of responsibility could rest on the simple idea that unless a man has the capacity and a fair opportunity or chance to adjust his behaviour to the law its penalties ought not to be applied to him.”); see also Winnie Chan & A.P. Simester, *Four Functions of Mens Rea*, 70 *CAMBRIDGE L.J.* 381, 388–93 (2011); John Gardner, *Wrongs and Faults*, in *APPRAISING STRICT LIABILITY* 51, 69–70 (A.P. Simester ed., 2005). These accounts tend to understand the value of mens rea in terms of the unfairness of punishing individuals absent notice rather than in terms of the ineffectiveness of attempting to guide individuals not to perform certain actions without adequately informing them of which actions are not to be performed. But regardless of why the function served by mens rea is important, this account, like mine, understands the criminal law to require mens rea in order to ensure that individuals are adequately informed of which actions the criminal law requires them not to perform.

<sup>100</sup> ROBINSON, *supra* note 80, at 125.

mens rea doctrines, in his view, perform the latter rather than the former function.<sup>101</sup> On Robinson's account, then, someone who brings about a particular prohibited result absent mens rea has failed to conform his behavior to the rules of conduct imposed by criminal law, but has done so in a blameless manner. By requiring the mens rea of knowledge, say, a criminalization rule does not instruct individuals they may act so long as they do not know that their action constitutes the actus reus of a crime; rather, the law forbids individuals from performing that actus reus *simpliciter* while holding some individuals blameless for failing to comply with the rule they have breached.

But this account completely misapprehends how the actual criminal law guides individual behavior. Consider, for example, someone considering whether to acquire property under circumstances indicating some risk that it has been stolen. What rule does the criminal law provide to govern his conduct? The Model Penal Code offense of "Receiving Stolen Goods" appears relevant: "A person is guilty of theft if he purposely receives, retains, or disposes of movable property of another knowing that it has been stolen, or believing that it has probably been stolen . . . ."<sup>102</sup> This rule convicts those who know or believe it probable that the goods to be acquired are stolen, and acquits them otherwise. Alternative mens rea provisions would produce different outcomes: if the law allowed only knowledge to suffice, those who merely believe the goods to be stolen would be acquitted, while if recklessness sufficed, the law would convict those who disregard a substantial risk that the goods received were stolen while believing they probably were not stolen. The defendant who merely perceives a substantial and unjustifiable risk that the goods he receives were stolen, then, would be acquitted under the Model Penal Code's offense definition but would be convicted and punished for an offense defined to require only recklessness instead.

Since conviction and punishment express censure and condemnation, only an offense defined in terms of recklessness would condemn and censure accepting goods

---

<sup>101</sup> See *id.* at 129–37. There is an exception: Robinson takes the mens rea provisions of inchoate offenses to serve the rule-articulation function. Since a substantial step toward commission of an offense, for example, may be entirely permissible absent any purpose to actually commit the offense, requiring defendants guilty of an attempted crime to possess the purpose actually to complete the attempt does specify the rule that individuals must follow. *Id.* at 133–35. Furthermore, Robinson takes some mens rea provisions to serve a third function of criminal law: grading different offenses as forms of wrongdoing. *Id.* at 132–33.

<sup>102</sup> MODEL PENAL CODE § 223.6(1) (AM. LAW INST. 1962).

if one recognizes a substantial and unjustifiable risk they were stolen: differences in the grade of mens rea employed produces differences in which actions the law condemns. And since punishment deters, those differences in mens rea also determine whether the threat of punishment provides individuals with a reason not to perform an action. It seems incredible that these differences could fail to influence how individuals are actually guided by the law. Someone who strives to comply with law and who recognizes a risk that the goods he may buy were stolen would surely be influenced in his deliberation by whether the law would in fact condemn and punish the act of buying them. And the mens rea provision used to define the offense would directly determine whether that action would be condemned and punished. Thus, it seems such mens rea provisions actually do have a genuine impact on how the law guides the conduct of those they bind.<sup>103</sup>

But because Robinson denies that mens rea provisions are relevant to rule articulation, he must claim that these possible offense definitions—which differ only in the mens rea they require—articulate identical rules of conduct. The rules disagree, in his view, about whether an individual who recognizes that goods might be stolen is blameworthy, and thus disagree about whether that individual is criminally liable, but even though only one authorizes convicting and punishing individuals who recognize a risk that the goods they accept might be stolen, they do not differ in whether they guide individuals to avoid such actions or not. Instead, these different possible offense definitions all promulgate a single rule of conduct simply prohibiting taking property that belongs to others.<sup>104</sup> Thus, he must hold that whether an individual would be punished for performing an action makes no difference to whether the law guides him not to perform it. But it is difficult to see how whether the criminal law permits an action could be unaffected in this way by whether it punishes that action. Indeed, standard accounts of how criminal law guides action insist that it guides action through the punishments it imposes: by punishing certain actions, it guides individuals not to perform those actions that it punishes. And, obviously, differences

---

<sup>103</sup> For a similar argument that mens rea must guide conduct, see R.A. Duff, *Rule-Violations and Wrongdoings*, in *CRIMINAL LAW THEORY: DOCTRINES OF THE GENERAL PART* 47, 69–70 (Stephen Shute & A.P. Simester eds., 2002).

<sup>104</sup> ROBINSON, *supra* note 80, at 215 (“You may not damage, take, use dispose of, or transfer another’s property without the other’s consent.”).

in the mens rea required for an offense produce differences in who will be punished for performing it. If the law can still guide individuals not to perform certain actions without imposing punishment for performing them, as Robinson's theory appears to require, it would seem unnecessary to impose punishment at all in order to guide conduct appropriately.

Furthermore, Robinson's own suggested rule of conduct, which does not incorporate mens rea provisions, clearly would not provide adequate guidance. How should an individual who recognizes that goods might be stolen apply a rule prohibiting receiving them just in case they are, in fact, stolen? That rule directs conduct based on the actual state of the world. As I have argued, however, individuals cannot deliberate based directly on the actual state of the world; instead, they deliberate based on the information about the world that they possess. To guide conduct effectively, rules must instruct individuals how to make the decisions that they actually face: an agent who must decide whether to accept goods that he recognizes might be stolen requires a rule specifying whether he may accept goods that he recognizes might be stolen. Offense definitions that incorporate mens rea provide that guidance: by specifying what mental states a defendant must possess in order for his action to be criminal, mens rea provisions guide action in terms of the information agents actually possess in deliberation. But Robinson's proposed rule of conduct fails to provide that guidance, for the criterion it provides to determine whether acting would be permissible depends on information individuals do not actually possess.

Due to human error, mistake, and accident, individuals cannot guarantee that their actions will never cause harm. Because compliance with that standard lies beyond human capacities, it is unreasonable for the law to demand it. But by punishing individuals simply because they cause a particular result, criminalization rules framed in terms of results alone make precisely that demand. Instead, individuals must appropriately balance the goal of avoiding certain severe harms with the various other goals that they might pursue. Thus, rather than mandating the unattainable avoidance of all harms, criminalization rules guide conduct effectively by mandating that individuals strike that balance appropriately, specifying the degree

of care that individuals must take in avoiding harmful results rather than simply instructing individuals to avoid them. For that reason, criminalization rules require mens rea, specifying that actions performed with particular mental states are sufficiently dangerous that they must be ruled out entirely, thereby leaving individuals free to choose among the actions that remain.

## 2. *Grading by Results*

Obviously, the decisive influence exerted by criminal punishment is essential to this argument for mens rea: criminalization rules must exempt some of an agent's options from the threat of criminal liability only because criminalization rules identify options that agents must exclude from deliberation, and rules cannot guide conduct at all if they require all options to be excluded. As I have noted, however, not all methods of guiding conduct by law involve this decisive force: it is possible to impose some negative consequences on those who perform an action without condemning those who perform it or aiming to fully deter them from doing so. Such remedies, which include taxes, regulatory fines, and arguably compensatory damages, provide individuals with a non-decisive rather than a decisive reason to avoid the conduct for which they may be liable. By imposing a negative consequence, the law creates a reason that influences individuals' deliberation, but it does not displace individuals' judgment entirely by dictating what options they may and may not choose. Instead, the law gives individuals a reason to prefer some options over others but ultimately defers to their own judgment about how to balance that reason against any others that might apply. Consequently, my argument concerning the importance of mens rea in criminal law does not apply to rules employing these remedies: they need not limit liability based on whether an individual possesses a particular mental state. And, indeed, the law incorporates no such requirement: the rules specifying the scope of individuals' liability to remedies of this sort often are framed in terms of results alone.<sup>105</sup>

---

<sup>105</sup> Thus, on my account even the criminal law may legitimately omit a mens rea requirement in a particular offense definition as long as the penalties imposed for committing that offense exert a non-decisive rather than decisive influence on deliberation. The Model Penal Code, for example, defines the limits of strict liability offenses in terms of the sanctions authorized for the commission of an offense: it allows strict liability only for "violations," MODEL PENAL CODE § 2.05(1)(a) (AM. LAW INST. 1962), which are defined as offenses that do not constitute crimes and for which the only authorized sanctions are fines, forfeitures, and other civil penalties, *id.* §



Grading rules, like these other legal remedies, aim to promulgate a non-decisive criterion: in particular, they guide individuals towards committing less serious offenses rather than more serious ones. Consequently, grading rules need not identify any binary classification among an agent's options. And since the function of mens rea requirements is to mark the division between options to be excluded from deliberation and options that are eligible to be chosen, grading rules need not employ mens rea. Instead, grading rules must only identify some basis for comparison among options—some criterion that individuals may employ to identify which of their options they should favor in deliberation and which they should not. And though a criterion employing results alone does not identify a basis for individuals to divide their options into two categories, it certainly does identify a basis for comparison among options: different options have different likelihoods of causing any result, thereby enabling individuals to distinguish options based on their likelihood of causing it. Grading offenses as more serious if they cause particular results conveys that actions causing that result are worse than actions that do not cause it, and punishing offenses that cause the result more severely gives individuals a reason to favor actions that do not cause it. Of course, individuals cannot apply these rules directly—they cannot choose an option if and only if that option will not in fact cause the more serious result. But they can apply these rules indirectly, preferring options less likely to cause that result. Of course, since this reason is non-decisive, it will not on its own determine which action individuals must choose; instead, individuals must act based on the balance of all reasons. But the additional reason provided by the grading rule will shift the balance of reasons towards less dangerous actions, thereby reducing in aggregate the frequency with which more serious results are expected to

---

1.04(5). Similarly, Mark Kelman has argued that strict liability may sometimes deter more efficiently in contexts such as these. His key example involves a strict liability violation of selling liquor to minors, enforced by a fine. Mark Kelman, *Strict Liability: 1. An Unorthodox View*, 4 ENCYCLOPEDIA OF CRIME AND JUSTICE 1512 1517 (Sanford H. Kadish ed., 1983). By imposing fines of this sort, though, the law does not identify a decisive criterion for individuals to employ in deliberation: fines do not plausibly express the kind of condemnation expressed by more severe sentences, and they are not sufficiently harsh that the threat of being fined will always outweigh whatever reasons a liquor store might have to pursue conduct that risks liability to a fine. Indeed, Kelman acknowledges that imposing strict liability for selling liquor to minors aims not to prohibit liquor stores from ever selling to minors but rather to induce them to take appropriate precautions against doing so. *Id.* This brings out the contrast with traditional criminal punishments: the prohibition and punishment of, say, intentional killing aims not to induce individuals to take appropriate precautions against intentionally killing but induce them never to intentionally kill (unless justified).

occur. Consequently, grading rules that grade by result can guide conduct effectively even though criminalization rules framed in terms of results alone cannot effectively guide conduct.

## **II. Grading Unintentional Risk-Creation**

Thus far, I have rejected the argument that because the criminal law's criminalization rules employ mens rea, its grading rules must employ mens rea too. But rejecting this argument obviously supports grading by results only negatively: it undermines one potential reason against grading by results, but it does not provide any reason in its favor. In this Part, then, I will present a positive argument for grading by results. That argument will not claim that grading by results necessarily constitutes a perfect method by which grading rules could incorporate the risk of unintentional consequences in their evaluations of the seriousness of criminal acts. Rather, in my view the law should grade offenses by their results primarily because there are no good alternatives: grading by results may not be an ideal method of grading, but it is the best method available. In particular, one obvious alternative exists to grading by results: perhaps the most natural way to guide offenders to reduce the risk that their actions will unintentionally cause harm would be simply to grade offenses directly by the risk they create. Nonetheless, this alternative will often fail: in many cases, grading by risk will estimate risk inaccurately and thus guide conduct ineffectively. To be sure, I do not claim that the law should grade by results alone, and indeed it does not do so: when grading by risk can succeed, it should be used. But because oftentimes it cannot succeed, the law should grade offenses by their results, as well.

In general, grading rules guide conduct by punishing certain offenses more severely than others. These differences in sentences express the judgment that the criminal acts punished more severely are more serious, and they provide individuals with a stronger incentive to avoid those acts. Thus, they guide individuals to avoid those offenses that the law punishes more severely. Grading by risk designates a family of grading rules rather than a single rule, since many different notions of risk and methods of measuring it could be employed in grading offenses. Despite these differences, though, all forms of grading by risk broadly guide individual behavior in

the same way. Each employs a particular method of estimating the risk created by each criminal act, which determines the sentence imposed for performing that act. Since these risk estimates determine the sentence to be imposed, grading by risk guides individuals to perform offenses whose risk is low according to the method of estimating risk that the rule employs. Consequently, the accuracy with which grading by risk estimates risk is crucial to whether it can successfully guide offenders' conduct. To defend grading by result, then, I will examine various methods of grading by risk, each of which employs a different approach to estimating the risk created by individual offenses. Some such approaches may grade accurately in some cases, and I agree that the law should grade by risk when it may do so accurately. But I will argue that in a wide range of circumstances each will systematically produce inaccurate risk estimates, and consequently will fail to effectively guide offenders to perform less risky offenses.

By contrast, a system of grading by results will generally guide offenders to reduce risk more effectively. The obstacles facing grading by risk arise from difficulties in making accurate estimates of risk and in communicating them to the offenders that must employ them in guiding their own conduct. But in grading by result, the law does not instruct offenders to reduce the risk created by their offenses, which would require the grading rule itself to articulate how offenders should measure risk, but rather merely instructs offenders which results to avoid. Offenders themselves, then, must translate that directive into action by relying on the best estimates of risk they themselves can make. Offenders' own risk estimates will not be perfect, and the law should grade by risks if it can identify, estimate, and communicate them more accurately than offenders can. But in many cases offenders can most effectively determine how to reduce the risk their own offenses create. Thus, requiring offenders to act based on their own risk estimates will reduce risk more effectively than requiring them to employ risk estimates made by government bodies and somehow communicated to offenders. Grading by result may not be ideal, but oftentimes it will be better than alternative grading methods. Though in some circumstances the law should grade offenses by risk, then, it should also grade offenses by their results.

In some respects, this argument is similar to a common philosophical defense of the role luck plays in judgments of moral blameworthiness. We often lack reliable evidence of the degree of risk created by actions that we evaluate morally. Absent that evidence, some philosophers argue, moral evaluation must rely on the results an action caused, which are our best evidence of the degree of risk that it created. We judge individuals to be more blameworthy when their actions cause harm, then, because that evidence indicates that their actions were in fact riskier.<sup>106</sup> My account is similar in that I take the importance of results in evaluating actions to arise from evaluators' systematic ignorance about the exact degree of risk that wrongful actions create. But I propose a different explanation for why it is reasonable to employ results in the absence of reliable evidence of risk. These philosophers' accounts agree that wrongdoing should be evaluated based on risk; they take results to be important only as evidence—the best available—of the risk an action actually did create.<sup>107</sup> My view instead takes more seriously the idea that evaluators cannot reliably estimate risk: if they are likely to estimate risk badly, the criminal law should not strive merely for them to estimate risk in the least bad way possible but rather should devise some other method of grading offenses that will accomplish its aims. Grading by result is such a method: it allows the law to guide individuals' behavior appropriately in the absence of reliable estimates of the risk created by any particular criminal act.

#### A. *Results, Proportionality, and Luck*

I will defend grading by results, then, by showing why each form of grading by risk would estimate risk inaccurately and thus guide conduct ineffectively. First, though, I will consider a broader objection to my methodology for defending grading by results. I have argued the law should grade offenses by their results because doing so would most effectively guide offenders' conduct. But it may seem that a different, more important kind of consideration decisively favors grading by risk, in which case a grading rule's effectiveness in guiding conduct might simply be irrelevant. In

---

<sup>106</sup> For approaches along these lines, see Nicholas Rescher, *Luck*, 64 PROC. & ADDRESSES AM. PHIL. ASSOC. 5, 15–17 (1990); Norvin Richards, *Luck and Desert*, 95 MIND 198, 200–02 (1986); and Judith Jarvis Thomson, *Morality and Bad Luck*, 20 METAPHILOSOPHY 203, 212–13 & n.5 (1989).

<sup>107</sup> As some critics have noted, it is not particularly plausible that results provide very good evidence of risk. *E.g.*, Lewis, *supra* note 17, at 57; Moore, *supra* note 8, at 250.

general, both ordinary moral intuitions and typical legal practice require criminal punishment to be proportional to the gravity of criminal wrongdoing: the sentence imposed for a particular offense must match that offense's seriousness.<sup>108</sup> And to many it seems that grading offenses by their results makes punishment proportional not to an offense's seriousness but rather to an offender's luck. I will argue that this consideration on its own cannot settle how the law should grade offenses, however. Grading by risk and by results each proportion punishment to the gravity of wrongdoing in some sense, and each allow punishment to depend on luck in some way. The law thus does not face a choice between either punishments proportional to the seriousness of an offense or punishments based on luck; rather, it faces a choice between different ways of measuring an offense's seriousness and different kinds of luck that might influence sentences. The importance of proportionality does not on its own constitute a reason to prefer one form of proportionality over another, then, just as the inappropriateness of sentencing according to luck does not constitute a reason to avoid one form of luck over another.

Whether a grading rule imposes punishments proportional to the gravity of wrongdoing will depend on what metric is employed to measure the gravity of particular criminal offenses. Since crimes breach binding rules of conduct—what makes an action a crime, that is, is that it satisfies the elements defining a particular kind of prohibited behavior—the law might naturally understand the gravity of a crime in terms of the magnitude by which it breached the relevant rule. In general, crime involves the improper exercise of rational agency: offenders disregard their reasons not to perform criminal acts, and their disobedience to the law's demands lies in this failure to respond appropriately to those reasons.<sup>109</sup> In defining and prohibiting offenses, the law requires individuals to treat the fact that an action satisfies the elements of a crime as a decisive reason not to perform it; by committing a crime, offenders fail to employ that reason in deliberation as the law demands, disregarding it rather than complying with it. Since offenses breach the law's demands by

---

<sup>108</sup> *E.g.*, 18 U.S.C. § 3553(a) (2018) (“The court, in determining the particular sentence to be imposed, shall consider— . . . (2) The need for the sentence imposed— (A) to reflect the seriousness of the offense . . .”).

<sup>109</sup> The failure to respond appropriately to one's reasons is a central element in many philosophical theories of blameworthiness, as well. *See, e.g.*, Matthew Talbert, *Moral Responsibility*, STAN. ENCYCLOPEDIA PHIL., 2.3 (Edward N. Zalta ed., Winter 2019 ed.), <https://plato.stanford.edu/archives/win2019/entries/moral-responsibility/>.

disregarding reasons that the law requires individuals to respect, the magnitude of their breach could be measured by the magnitude of the offender's failure to respond appropriately to those reasons. An offender misvalues his reasons: by offending, he implicitly judges his reasons against offending to be outweighed by reasons in favor, even though the law demands that he assign the reasons against acting a value that would outweigh any reasons in favor. The greater the divergence between his own improper valuation of those reasons and the valuation the law demands, then, the graver his wrongdoing.<sup>110</sup> And that divergence, of course, will be greater the stronger the disregarded reasons not to perform a particular crime are. Since offenders surely have stronger reasons to avoid riskier crimes than less risky ones, and since grading by risk sentences offenses based on the risk they create, it seems to punish crimes with severity proportional to their seriousness.

But the strength of the reasons that an individual disregards in committing a particular crime do not depend on whether it actually causes harm. To be sure, in some sense stronger reasons exist not to perform an action if it actually causes harm, for only if an action actually would cause harm can an agent avoid the harm by avoiding the action. But reasons of this sort are not available to individuals in deliberating about what to do. Individuals deliberate based on only those reasons of which they are aware; information that they lack cannot provide reasons to guide their deliberation. Thus, though individuals may employ their beliefs about an action in deliberating over whether to perform it, they cannot deliberate based on the actual results it will cause, for they have no access to that information when deliberating. If one action is likelier to cause harm than another, then in deliberation individuals have stronger reason to choose the latter, even on a particular occasion when the latter, not the former, actually will cause harm.

Consequently, grading by results seems to impose punishments that are disproportionate to the gravity of the offense punished, since it sentences offenders based on whether their offenses caused certain harmful results, even though the

---

<sup>110</sup> E.g., Gideon Yaffe, *Legal Reasons, Legal Desert, Legal Culpability: Reply to Guerrero, Kelly and Mendlow*, 24 J. ETHICS 295, 295 (2020) ("The greater the gap between [1] the reason-giving weight one is disposed to grant to a given fact in one's deliberations, and [2] the weight of the legal reason provided by the fact, the greater the criminal culpability for conduct manifesting that disposition.").

occurrence of those results does not change the reasons they disregarded in acting and thus does not change the degree to which their conduct deviated from the law's demands. While grading by risk proportions sentences to the gravity of wrongdoing, then, grading by result appears instead to proportion sentences to offenders' luck: harsher punishments are imposed not on those who fail to a greater degree to respond appropriately to the reasons governing their conduct but rather on those whose equally wrongful actions unluckily cause greater harm.<sup>111</sup> And since surely an offender's bad luck cannot justify punishing him more severely, the requirement for punishments to be proportional to the seriousness of the offense might seem on its own to require grading by risk rather than by result.

I will dispute neither that the severity of punishment must be proportionate to the gravity of wrongdoing nor that bad luck does not itself justify harsher punishment. But, in my view, neither of these principles on its own favors grading by risk rather than by result. First, grading by risk and by result do both impose punishments proportional to the gravity of the offense for which they are imposed; they differ, instead, in how they evaluate the gravity of particular criminal acts. There is one natural sense in which the degree of wrongfulness of criminal conduct lies in the degree to which the offender disregards the reasons that ought to govern his choice. But that is not the only way to understand the gravity of criminal wrongdoing. Though an action's actual results may not affect the reasons available to an agent deliberating whether to perform it, those results do affect something else about it—namely, the harm it ultimately causes. An arson that causes death kills someone, unlike one that does not cause death, and there is an obvious sense in which killing someone is worse than not killing anyone: the criminal law prohibits wrongdoing in

---

<sup>111</sup> On this analysis, then, the legitimacy of grading by result would be a legal analogue to the philosophical problem of moral luck, which broadly concerns whether luck may legitimately influence our moral assessment of actions. For the pair of papers that launched the modern debate concerning moral luck, see T. Nagel, *Moral Luck*, 50 PROC. ARISTOTELIAN SOC'Y 137 (1976); and B.A.O. Williams, *Moral Luck*, 50 PROC. ARISTOTELIAN SOC'Y 115 (1976). For further discussion of the topic, see generally the papers in MORAL LUCK, (Daniel Statman ed., 1993). Scholarly analysis of grading by results has frequently been framed—both by those who accept and those who reject it—as addressing the problem of whether luck may legitimately influence the criminal sentences imposed on offenders. See, e.g., Ashworth, *supra* note 3; Duff, *supra* note 9; R.A. Duff, *Whose Luck Is It Anyway?*, in CRIMINAL LIABILITY FOR NON-AGGRESSIVE DEATH 61 (C.M.V. Clarkson & Sally Cunningham eds., 2008); Kadish, *supra* note 2; Kessler, *supra* note 3; Lewis, *supra* note 17; Ken Levy, *The Solution to the Problem of Outcome Luck: Why Harm Is Just as Punishable as the Wrongful Action That Causes It*, 24 LAW & PHIL. 263 (2005); Smith, *supra* note 3; Steven Sverdlik, *Crime and Moral Luck*, 25 AM. PHIL. Q. 79 (1988). For a theorist who declines to equate the problem of grading by results and the problem of luck, though, see Moore, *supra* note 8, at 253–58.

order to protect certain interests, and an offense that causes death invades those interests more substantially than an offense that does not cause death.

Just as the gravity of criminal wrongdoing might plausibly be understood in terms of how an offender's exercise of rational agency deviated from the law's demands, then, it also might plausibly be understood in terms of how an offense invaded the interests the law aims to protect. On this interpretation, the greater the degree of that invasion, the more serious the crime.<sup>112</sup> And because the law aims to protect persons and their things, the results of a crime do affect whether it invades those interests, and thus do affect its seriousness as an offense. In this sense, results do affect the gravity of wrongdoing, and grading by results does proportion the severity of punishment to the gravity of wrongdoing, so understood. On its own, then, the requirement of proportionality favors neither grading by risk nor grading by result: both forms of grading are proportional, but in different ways. The relevant question is not whether punishment must be proportional at all but rather what it must be proportional to.

Similarly, the dependence of results on chance cannot alone establish that the scope of criminal liability cannot depend on results. For it is emphatically untrue that some general principle of law precludes individuals' legal liability from depending on chance. Chance is ubiquitous in the law. It dramatically affects the existence and scope of individuals' legal liability across a wide range of doctrinal areas: the lucky investor whose investments succeed owes income tax while the unlucky investor whose equally promising investments fail receives a deduction against other income; the lucky insurance company whose insureds avoid accidents faces little contractual liability while the unlucky insurance company whose insureds suffer accidental injury may owe considerable sums; the lucky negligent driver whose victim survives with minor injuries faces lower tort liability than the equally negligent unlucky driver whose victim is badly injured or killed. Luck affects liability even within criminal law: the unlucky heir who kills his long-lived childless aunt to secure access to his inheritance is guilty of murder, for example, while the lucky heir whose aunt died by

---

<sup>112</sup> For an argument that justifies grading attempts by their results using considerations of this sort, see DUFF, *supra* note 9, at 190–92; and Duff, *supra* note 9, at 36–37.



chance years ago commits no crime and faces no liability. Indeed, even grading by risk allows liability to depend on chance: an unlucky offender may create an elevated risk due to some fortuitous circumstance—the presence of a homeowner during a burglary, or the medical condition of a robbery victim, or high winds during an arson. But none of these ways in which chance determines liability seems objectionable, and it is hardly plausible that all of them must be purged from the law. Because chance is ubiquitous, legal liability will depend on luck to some extent; whether grading by results is justified depends on which forms of chance may legitimately affect liability and which may not.

Furthermore, although these forms of legal liability do all depend on chance, none of them treats a defendant's bad luck as a positive justification for imposing liability. Tax liability depends on a taxpayer's income; contractual liability on the promisee's expectation; tort liability on the victim's injuries. Of course, chance may influence a taxpayer's income, or a promisee's expectation, or a victim's injuries. But the greater liability of the unlucky is not justified because they are unlucky. Rather, the particular basis the law uses to determine liability is justified on its own terms: some reason justifies taxing individuals by their income, holding promisors to their promisees' expectations, or requiring injurers to compensate their victims. This underlying justification for the relevant principle of liability justifies the liability imposed on particular defendants: if injurers may justifiably be required to fully compensate victims, say, then those who cause worse injuries must automatically pay more. If that basis for liability is justified, no further justification for the role of chance is needed: if, say, requiring injurers to fully compensate their victims for their injuries is justified in general, then since injuries depend on chance, it follows automatically that tort liability may justifiably depend on chance. There is no separate problem of luck, conceptualized as such. Whether liability may justifiably depend on some form of chance depends on whether the basis the law employs to determine liability is justified: luck may affect liability to the extent that it affects whatever basis legitimately determines liability, and it may not otherwise. Whether the law ought to grade by risk or by result, in either case chance may legitimately influence criminal

liability by influencing either risk or result, respectively. The only question is what basis for liability the law ought to employ.

How, then, should the law decide whether to employ risk or results? One possible approach—appealing to ordinary moral intuitions about just punishment—is unlikely to resolve this question, because such intuitions appear torn between each option. On the one hand, in the abstract the sort of luck involved in grading by results strikes many as intuitively objectionable. As J.C. Smith puts it, “it might seem obvious that when we have established all the facts, a person’s liability should not depend on matters of chance”; he suggests it is “shocking in the second half of the twentieth century” that this principle is not always satisfied.<sup>113</sup> And arguments against grading by result almost universally run through hypotheticals in which two otherwise indistinguishable offenders produce different results because of some factor obviously outside their control—a gust of wind, a firearms malfunction, the flight of a bird—to elicit the intuition that an offender’s sentence should not depend on a difference of this sort.<sup>114</sup> But although in the abstract it seems intuitively wrong for individuals’ sentences to depend on factors outside their control, it also intuitively seems that those who cause harmful results should be punished more severely than those who do not. Grading by results is ubiquitous—Sanford Kadish describes its “near universal acceptance in Western law”<sup>115</sup>—and it is consistent with most individuals’ ordinary intuitions about interpersonal morality and law.<sup>116</sup> While it is

---

<sup>113</sup> Smith, *supra* note 3, at 63; see also Feinberg, *supra* note 3, at 118 (“In these examples the moral blameworthiness of criminals is identical, yet the punishment is much more severe in the one case than in the other. Unless there is some reasonable explanation for this discrepancy, the sentences seem to be more arbitrary than rational, the difference between the fates of A1 and A2 being determined not by their deserts but by luck, plain and simple.”); H.L.A. HART, *Intention and Punishment*, in PUNISHMENT AND RESPONSIBILITY: ESSAYS IN THE PHILOSOPHY OF LAW, *supra* note 99, at 131 (“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?”); Kadish, *supra* note 2, at 691 (“It would be the same if we allowed every equally guilty offender a throw of the dice—a throw of six or less and we halve the punishment. The two offenders end up being punished differently even though they are identical in every non-arbitrary sense.”); Lewis, *supra* note 17, at 67 (“Most likely [our present practice] isn’t [just], but I don’t understand why not.”).

<sup>114</sup> ALEXANDER, FERZAN, AND MORSE, *supra* note 3 at 172; Alexander, *supra* note 3, at 8–10; Ashworth, *supra* note 3, at 108–09; Feinberg, *supra* note 3, at 117–18; Kadish, *supra* note 2, at 688; Schulhofer, *supra* note 3, at 1498; Sverdlik, *supra* note 111, at 79.

<sup>115</sup> Kadish, *supra* note 2, at 679; see also Schulhofer, *supra* note 3, at 1498–1503.

<sup>116</sup> DUFF, *supra* note 9, at 189 (“[W]hat my relief expresses is the thought that to have a failed attempt at murder on my conscience is quite different from having an actual murder on my conscience . . . .”); Kadish, *supra* note 2, at 688 (“[M]ost of us do in fact make judgments precisely of this kind. Doesn’t it seem natural for a parent to want to punish her child more for spilling his milk than for almost spilling it, more for running the family car into a wall than for almost doing so?”); PAUL H. ROBINSON & JOHN M. DARLEY, JUSTICE, LIABILITY, AND BLAME: COMMUNITY VIEWS AND THE CRIMINAL LAW 169–81 (1995) (surveying community views on the felony murder ruse). For a

intuitive that factors outside a defendant's control should not influence the severity of his sentence, then, it is also intuitive that the results of his offense should influence the severity of his sentence. And, of course, those results often depend on factors outside his control. Thus, moral intuition alone seems to deadlock between grading by risk and grading by result.

I propose instead to resolve the dispute between grading by risk and by results by identifying which basis of liability better achieves the criminal law's purposes in imposing criminal liability. Either form of grading, I have noted, produces proportional punishment; they differ, instead, only in whether they designate risk or results as the basis to which punishment must be proportional. But both forms of grading, on my analysis, serve the same function: each aims to guide conduct by providing offenders with a criterion to employ in choosing the manner in which they perform offenses. If grading rules aim to provide a particular kind of guidance, then at the very least they ought to do so effectively. But grading by risk will produce inaccurate risk estimates, and therefore it will not effectively guide offenders to perform less risky offenses. The rules of criminal law would obviously be counterproductive if they fail to achieve their own purposes, whatever those purposes may be. Thus, because grading by results rather than by risk better enables grading rules to achieve their purposes, the law should grade offenses by their results. This principle determines the manner in which punishment must be proportional: the seriousness of an offense must determine the severity of the sentence imposed for performing it, and the seriousness of an offense must be measured in part in terms of results, not only risk. And because results are a legitimate basis with which to determine the scope of criminal liability, the dependence of results on chance does not allow luck to exert any unjustified influence on the scope of criminal liability. Luck cannot itself justify imposing liability, but because any basis for determining criminal sentences may be influenced by chance, the influence of luck on results is not on its own a bar to the influence of results on liability. If for other reasons results are a justified basis for determining the extent of an offender's criminal liability, then

---

defense of grading by results based primarily on these intuitive considerations, see Moore, *supra* note 8, at 263–71.

the influence chance may have on results—like the influence chance exerts on any basis of legal liability—does not preclude results from legitimately influencing the extent of liability.

This argument in defense of grading by results appeals to the particular function I have ascribed to grading rules: since those rules aim to guide offenders to perform less risky offenses, they should assign criminal liability in a manner that does effectively guide offenders to perform less dangerous offenses. But my critique of grading by risk does not depend on the claim that grading rules serve this particular function. Grading rules that grade by risk employ a distinctive strategy to pursue their objectives—namely, they proportion sentences to risk. If they cannot accurately estimate the risks crimes create, they cannot achieve their objectives, whatever those objectives may be. Grading by risk fails on its own terms, because it will systematically fail to accurately estimate the risk that offenses create. Thus, though I will explain how inaccurate risk estimates prevent grading by risk from effectively performing the function I ascribe to grading rules, the scope of my critique is broader. The law does not genuinely face a choice between grading by results and grading by risk, because in most circumstances it cannot accurately grade by risk at all; instead, it must either grade by results or try and fail to grade by risk. If offenses could be graded by risk accurately, doing so might well be preferable to grading by result, and to the extent that the law can accurately estimate the risk offenses create I agree that it should grade offenses according to those estimates. But whatever theoretical advantages there might be to grading offenses by their actual risk cannot justify sentencing offenders according to inaccurate estimates of the risk their offenses create: a system that tries to impose proportional punishments but that cannot accurately measure the basis of liability simply fails to impose punishments that are actually proportional. Of course, perhaps grading by risk would still be justified were risk the only possible basis for assigning criminal liability. But it is not, since grading rules may employ results, instead. Because the law cannot estimate risk accurately, it should employ a method of grading offenses that does not require it to do so.

### *B. Risk Estimates Through Law*

Any grading rule that grades by risk must somehow estimate risk in sentencing individual offenders. By determining offenders' sentences, these risk estimates play a central role in guiding offenders' conduct, for the law communicates its judgments about the seriousness of different offenses through the sentences it imposes. How, then, might the law estimate the risk created by particular criminal acts? In general, the risk an action creates depends on two kinds of factors: particular facts about the nature of an action and the circumstances in which it was performed, and general rules that state how those particular facts affect the likelihood that the action will cause certain results. The risk that a particular robbery will cause a person's death, for example, depends both on whether the robber was armed and on the likelihood that armed robbery will cause death. To estimate the risk created by particular criminal acts, then, any system of grading by risk must identify both factual details about each offense and probabilistic rules specifying how those factual details affect the risk each offense creates.

How might the law identify each of these two components that determine the risk created by particular criminal acts? In general, criminal law involves two characteristic types of legal proceedings: some lawmaking body writes criminal laws, then some adjudicative body applies those laws to particular defendants. Grading by risk, then, could estimate risk using either of these procedures. Adjudication could identify either of the two kinds of information upon which risk estimates depend: courts could identify particular facts about the manner in which a crime was committed, and they could determine how those facts affected that crime's likelihood of causing a particular harmful result. By contrast, legislatures cannot plausibly identify particular facts about the manner in which any particular criminal act was committed. Lawmaking is general and prospective: it articulates rules that apply to entire categories of actions occurring after the rule's enactment. Facts about particular actions, however, are particular and must be identified retrospectively: they concern only one single act rather than applying across an entire category of actions, and only after a crime actually has been committed can anyone identify how it was committed. Thus, grading rules cannot plausibly task legislatures with identifying how particular

criminal acts were committed. But they could require legislatures to identify the general and prospective rules that govern how particular facts affect the risk an offense creates.

Consequently, there are two strategies that a system of grading rules might employ in order to grade offenses by their risk. Each requires courts to identify the manner in which each criminal act was performed, while they differ in whom they require to estimate the risk criminal acts create given how they are performed: on one approach, both these tasks are assigned to the courts, whereas on the other approach courts determine how a crime was performed but defer to the legislature's judgments as to the risk it thereby created. The former approach, which requires courts to estimate risk based directly on the facts of each crime, grades offenses directly by risk: its grading rules sentences offenders based directly on the risk their offenses created, thereby assigning the task of estimating risk wholly to the courts that impose sentences. The latter approach requires the legislature to promulgate rules estimating the risk created by different kinds of offenses; based on their own investigation of the particular facts involved in how each offense was performed, courts would then apply those rules to estimate the risk each offense created.<sup>117</sup> This approach grades offenses by risk factor: prospective grading rules identify which kinds of offense are more or less risky, and courts sentence offenders based on how those rules estimate the risk that each offense created.

I will argue that neither grading by risk factor nor grading directly by risk will accurately estimate the risks created by particular offenses. The law's risk estimates depend on both how it identifies facts about each offense and how it employs those facts to determine the likelihood of harm; its risk estimates will be accurate only if it accurately estimates each of these two factors. Different methods of grading by risk differ in which entity they make responsible for each part of risk estimation, but because both grading by risk factor and grading directly by risk will carry out each part inaccurately, both will inaccurately estimate the risks individual offenses create.

---

<sup>117</sup> In theory, the legislature could promulgate rules for estimating the risk offenses create and separate rules for how to sentence offenders based on the risk their offenses created. In practice, though, it will ordinarily be simpler to rely on a single set of rules that directly specifies the sentence to be imposed for performing offenses in various ways, based on the legislature's judgments as to the degree of risk created by performing offenses in those ways.

First, on either approach, courts must identify the facts about each offense that will be employed to estimate its risk. An enormous range of details can influence the risk any given offense creates, however, many of which will be obscure or highly specific. Though courts may accurately identify some such details, they surely will not identify all, and therefore the law's risk estimates will be based on incomplete information about each offense. Second, either courts or legislatures could employ those facts to estimate the risk created by each offense. But although each of those entities may have some capacity to estimate risks accurately, it is extremely difficult to precisely quantify the exact likelihood that any single action will cause any particular result. Thus, both methods of grading by risk will frequently fail to correctly estimate the likelihood that an offense performed in the manner found by the court would produce a particular result. Thus, grading by risk will first incorrectly identify how particular offenses were performed, then incorrectly identify the likelihood that an offense performed in that manner will cause harm. Because of these two sources of inaccuracy, the sentences imposed on offenders under grading by risk will not accurately track the risks their offenses created.

As a result, neither form of grading by risk will guide conduct effectively. Because grading rules guide offenders to avoid offenses that are punished more severely, an offender successfully influenced by grading rules will deliberate and choose how to offend based on how he believes various possible offenses would be punished by law. And since under grading by risk the law's risk estimates determine the sentence imposed for an offense, offenders will deliberate based on how they believe the law would estimate the risk created by each offense. If the law estimates risk inaccurately, then, it will guide offenders to perform offenses that are less risky according to the law's estimates, even if in fact they are in fact riskier. If courts cannot identify some particular fact about the manner in which an offense was performed, then even if that fact would affect the risk an offense creates, grading by risk cannot guide offenders not to offend in that manner, for courts will estimate risk identically and thus sentence offenders identically regardless of whether they do. Similarly, if courts or legislatures cannot accurately identify and communicate to offenders which offenses are riskier than others, they cannot guide offenders to

perform the less risky ones. Thus, because grading by risk will fail to accurately estimate the risk offenses create, it will fail to effectively guide offenders to perform less risky offenses.

My arguments against grading by risk, then, center on the structural obstacles that generally confront the law in producing estimates of the risks individual actions create that may be communicated to offenders and may guide their deliberations about how to offend. But although the existence of these structural obstacles makes grading accurately by risk difficult, it is not impossible. In some circumstances, those obstacles will not exist or may be overcome, and my arguments do not identify any reason not to grade by risk in those circumstances. I fully accept this conclusion. The problem with grading by risk, in my view, is not that it can never succeed at guiding offenders to perform less risky offenses but rather merely that it cannot always succeed at doing so. The law should grade by risk, then, when the conditions doing so effectively may be met. But because of the obstacles that confront grading by risk, some risks cannot be accurately incorporated into the sentences grading by risk imposes and therefore cannot be effectively reduced through grading rules that grade by risk. To be sure, I will not give a detailed account of exactly when grading by risk will and will not succeed; my goal here is to articulate the theoretical considerations that govern grading, not to produce a detailed blueprint for sentencing law. Because of the obstacles that confront grading by risk, some combination of grading by risk and by result will guide conduct most effectively, but I will not try to determine exactly what role should be given to each. Rather, I aim merely to defend some role for results in grading by identifying limits on the effectiveness of grading by risk.

### *1. Grading by Risk Factor*

Grading rules that grade by risk factor divide the task of estimating risk between the entity that makes criminal law, such as the legislature or a sentencing commission, and the courts that apply it. First, lawmakers write sentencing rules that reflect their judgments as to which kinds of offenses are riskier than others. Such rules each define a class of offenses, each of which possesses some feature that increases its likelihood of causing harm, then impose a higher sentence on all offenses in that class, commensurate with the increased risk they create. Courts apply these



rules when sentencing individual defendants, analyzing each offense to identify whether it possesses the features that would require the application of various sentencing rules, then imposing the sentence required by whichever rules apply to the offense.

Most systems of sentencing law grade by risk factor to some extent simply in defining offenses. Those definitions must specify the default sentences to be imposed for committing each offense. In defining a crime, the law defines a category of criminal acts characterized by certain features—namely, the statutory elements that define the offense. Ordinarily, an offense causes some harm simply by satisfying those elements, and the default sentence for the offense should depend in part on the magnitude of that harm. But in addition, some risk often exists that an action satisfying those elements will also cause additional, unintended harms, and that risk should also affect the sentence imposed. Thus, the elements that define an offense constitute a kind of risk factor. A system of grading by risk factor might employ only grading rules of this sort—that is, every offender who commits the same crime, as defined by statute, might receive the same sentence.<sup>118</sup> Alternatively, grading rules could define risk factors that distinguish among different ways of performing the same offense. That is, if the statutory definition of an offense is broad enough to encompass conduct varying considerably in dangerousness, lawmakers could promulgate additional grading rules that distinguish in sentencing between versions of that offense that are more or less dangerous.<sup>119</sup> Though these different versions of grading by risk may differ considerably in the detail with which they define risk factors, they share a common approach to estimating the risk created by each offense.

Under grading by risk, then, courts determine the manner in which an offense was performed, and lawmakers determine the risk created by different offenses performed in different ways. Sometimes, certainly, each will be able to discharge its

---

<sup>118</sup> Certain Model Penal Code provisions, for example, define particular offenses as always constituting a felony of the same degree. *E.g.*, MODEL PENAL CODE § 230.2 (AM. LAW INST. 1962) (defining incest as a third-degree felony). Because the Model Penal Code allows judges some discretion in sentencing, though, different offenders might be sentenced to different terms even if each commits a felony of the same degree. *See id.* § 6.06.

<sup>119</sup> Such rules appear in both the Model Penal Code and the Federal Sentencing Guidelines. *See, e.g.*, U.S. SENTENCING GUIDELINES MANUAL § 2B3.1(b)(2) (U.S. SENTENCING COMM’N 2021) (grading robbery as more serious if it involves a firearm, other dangerous weapon, or threat of death); MODEL PENAL CODE § 221.1(2)(b) (grading burglary as a second-degree felony if the offender “is armed with explosives or a deadly weapon”).

task accurately. First, as illustrated by existing examples of grading by risk factor, lawmakers do seem able to identify some factors that increase the risk that an offense will cause harm. The common practice of grading violent crimes as more serious if they involve dangerous weapons, for example, reflects the judgment that offenses involving such weapons are likelier to cause death or injury,<sup>120</sup> just as the practice of grading forgery based on what document was forged reflects the judgment that different sorts of forged documents are generally likely to cause different amounts of harm.<sup>121</sup> And while such rules may sometimes inaccurately quantify the precise degree to which some risk factor increases the risk an offense creates, oftentimes lawmakers will identify risk factors more reliably than offenders themselves. The individual forger, say, may have relatively little insight into the more attenuated consequences of a particular act of forgery, and thus he may not appreciate how different kinds of forgery create greater or lesser risks. By contrast, because lawmakers have considerable resources for empirical investigations and extensive expertise in identifying and proscribing dangerous conduct, sometimes they plausibly will be able to identify features of offenses that increase the risk they create but that offenders themselves would fail to recognize.

Second, the court that sentences an offender is often able to determine which risk factors characterized his offense. In order to adjudicate a defendant's guilt, a court ordinarily must determine what kind of action he performed—in particular, whether the action satisfied the elements of the offenses charged in the indictment. Many risk factors involve similarly general facts about the type of action a defendant performed—whether, say, he forged one type of document rather than another, or whether he used a dangerous weapon to commit a violent crime. While obviously no method of factfinding is guaranteed never to err, the evidence typically available to courts will often be adequate to establish these facts about the general manner in which an offense was performed. Witnesses may testify about whether or not a robber held a gun, for example, or a forged document may itself be presented in court.

---

<sup>120</sup> See *supra* note 119.

<sup>121</sup> E.g., U.S. SENTENCING GUIDELINES MANUAL § 2B5.1(b)(2) (defining a separate guideline, with a higher base offense level, covering counterfeit bearer obligations of the United States); MODEL PENAL CODE § 244.1(2) (grading forgeries according to the type of document forged).

Grading rules like these define risk factors that courts can identify reasonably accurately. Consequently, at least sometimes grading rules that grade by risk factor do produce reasonably accurate estimates of risk. If courts can correctly identify certain characteristics of individual offenses, and if legislatures can correctly quantify the degree to which those characteristics affect the risk those offenses create, then imposing sentences in part based on those risk factors will partially proportion the severity of criminal punishment to the risk that criminal offenses create.

When these conditions are met, grading by risk factor effectively guides offenders to perform less risky offenses. By identifying risk factors in advance, grading rules that grade by risk factor communicate which kinds of offenses courts will estimate to be more or less risky, thereby influencing offenders to perform offenses estimated to be less risky. Sometimes, most offenders will likely know already that certain offenses are riskier—surely, for example, few need to be informed that an armed robbery is likelier to cause death than an unarmed one. By punishing riskier offenses more severely, then, grading rules that grade by risk factor guide offenders to actually avoid those risk factors in the offenses they perform. In other cases, though, grading rules may guide offenders to avoid risks that they themselves would not have recognized on their own. For while some risk factors may be obvious, others may not—perhaps, for example, some offenders may not recognize that forging an issue of stock creates a greater risk of harm than forging a deed.<sup>122</sup> In these cases, grading rules actually inform offenders which offenses are riskier in addition to guiding them to perform the less risky ones. By employing grading rules that grade by risk factor that do accurately estimate the risk offenses create, then, the law can effectively guide offenders to perform less risky offenses; indeed, such grading rules can sometimes influence offenders to reduce risk more effectively than they could on their own. As I have noted, in my view the law ought to employ some combination of grading by risk and by results rather than grading by results alone: though limitations on the accuracy of the law’s risk estimates prevent grading by risk from always guiding conduct effectively, I agree that the law should grade by risk

---

<sup>122</sup> See MODEL PENAL CODE § 244.1(2). Of course, I assume here that the purpose of this grading rule is risk-reduction, which may be incorrect.

when it can do so accurately. Thus, if lawmakers have identified a factor that consistently increases the risk an offense creates, and if courts can reliably determine whether any particular offense features that risk factor, the law should employ a grading rule that sentences offenders more severely for offenses characterized by that risk factor.

Nonetheless, oftentimes at least one of those conditions will not be met. Grading by risk factor can accurately estimate risk only if courts can accurately determine which offenses feature which risk factors. Because courts can easily identify, say, whether an offender employed a gun in performing his offense, they can accurately apply rules that grade offenses by whether they were performed with a gun. But countless details about the particular manner in which an offense was performed may affect its likelihood of causing any particular harmful result, and not all such details will involve facts as readily apparent as the presence of a weapon or the kind of document that was forged. The precise degree of force applied in a physical confrontation, the precise direction in which a gun was fired, the precise angle at which a car sped around a turn—all these factors may affect exactly how dangerous an action was. And it is quite implausible that courts can consistently identify all such factors accurately.

Most often, courts gather information through observation, including witnesses' observations of the crime itself, investigators' observations of the crime scene, and the court's own review of objects or documents entered as evidence at trial. Certainly, observations of this sort can reliably identify general facts about what kind of action a defendant performed. But ordinary observation has only a limited ability to draw fine distinctions between the precise details of each offense, even though many such precise distinctions do affect the degree of risk an offense created. Thus, while a witness's ordinary observations can easily identify whether an offender used a gun or whether he fired it, and the court's own observations can identify the kind of document a defendant forged, ordinary observations alone will often mistake the precise degree of force used in assault, say, or the precise location on a victim's body at which a gunshot was aimed. And this kind of observational evidence is the best evidence that courts are likely to have. In controlled settings, of course,

experimental techniques can often produce exact measurements, but only rarely will some technical measurement apparatus already be in place when a crime is committed. If observational evidence is unavailable, courts will most likely have access only to circumstantial evidence. And while certainly circumstantial evidence will yield some information about the manner in which an offense was performed, it is likely to be even less precise than observational evidence. Finding stolen goods or a weapon in a defendant's apartment or finding chemical residue on his clothes or belongings, say, is good evidence that he committed a crime, but evidence of that sort can hardly be used to determine the exact manner in which the crime was committed. Thus, courts will simply lack reliable access to the evidence that would be required to accurately determine whether many risk factors characterize individual offenses. And since courts cannot accurately identify these risk factors, courts grading by risk factor will produce risk estimates that inaccurately reflect the risk those risk factors contribute.

Just as courts cannot identify every aspect of an offense that affects the risk it creates, lawmakers similarly cannot promulgate rules that define any factor that could affect the risk an offense creates. In general, the entities and individuals that apply a grading rule differ from those that produce it: lawmakers write the grading rules that specify how offenders should be sentenced, but courts and offenders themselves must apply those rules either in imposing sentences or in guiding their own deliberation and choice. Grading rules must communicate with courts and offenders: in promulgating any grading rule lawmakers must describe linguistically which offenses the rule judges to be more serious and thus punishes more severely. A grading rule that grades by risk factor, then, must include an explicit, verbal definition of that risk factor, which must describe it clearly enough to be communicated to and understood by the courts and offenders that must apply it. But there is no reason to suppose that every aspect of every offense that affects the risk it creates can be given a clear linguistic definition that all others may readily understand. In general, humans' ability to differentiate between experiences outruns our ability to comprehensibly describe those differences to others in language. Individuals can distinguish between different ways of performing an action—ways of physically assaulting someone that differ in

their likelihood of causing injury, for example, ways of driving a car that differ in their likelihood of causing an accident, or ways of robbing a bank that differ in their likelihood of causing a violent confrontation—even if they cannot always describe that difference so as to allow another speaker of the same language to grasp it or to apply it to their own experience.<sup>123</sup> These differences between different sorts of behavior may mark differences in the risk those actions create, but if there is no way to mark that difference in language, no grading rule could define a risk factor that captures how these aspects of criminal conduct affect how risky an act is.

Second, grading rules that grade by risk factor require courts not simply to identify how a particular aspect of some action affects the risk created by that action but rather to define rules that quantify the universal impact each risk factor produces on the degree of risk created by any offense that features it. When employing a particular risk factor to estimate the risk that one single action creates, we ordinarily need to consider the likely causal effects of that risk factor only given all of that action's other actual features: even if we do not know how that risk factor would affect risk in very different contexts, we can nonetheless estimate how it affects the risk created by the particular offense in question so long as we understand its effects in that context. But grading rules that grade by risk factor do not express judgments

---

<sup>123</sup> Philosophers have explained why language cannot always describe or communicate distinctions in experience by appealing to various aspects in which it differs from the concepts employed in language, such as its fine-grained or analogue nature, *see* CHRISTOPHER PEACOCKE, *A STUDY OF CONCEPTS* 67–69 (1992), its richness, *see* Richard G. Heck, Jr., *Nonconceptual Content and the "Space of Reasons,"* 109 *PHIL. REV.* 483, 489–90 (2000), or its context-dependence, *see* Sean D. Kelly, *The Non-Conceptual Content of Perceptual Experience: Situation Dependence and Fineness of Grain,* 62 *PHIL. & PHENOMENOLOGICAL RESEARCH* 601, 606–08 (2001). Philosophers have often analyzed the inability of language to fully capture the content of experience in order to argue that the contents of experience are non-conceptual—in particular, to argue that individuals may have an experience of a certain type without being able to conceptually specify the contents of that experience, for example by describing it in utterances. *See, e.g.,* Heck, Jr., *supra*, at 484–89. Obviously, this position has some similarities with the claim I advance here. But my claim is considerably weaker: I do not deny that individuals can specify the contents of experience conceptually but rather deny only that lawmakers can specify the contents of a class of experiences, in advance, in a manner that will enable individuals to consistently recognize those experiences when they have them. Thus, my claim does not depend on nonconceptualism. For example, John McDowell has defended conceptualism by arguing that the contents of experiences can be articulated by using demonstrative concepts, like “this shape” or “this color,” or by using concepts tied to an individual’s capacities to recognize similarities between current and previous experiences. *See* JOHN MCDOWELL, *MIND AND WORLD* 56–60 (1994); Christopher Peacocke, *Does Perception Have a Nonconceptual Content?*, 98 *J. PHIL.* 239, 244–52 (2001). But, obviously, the criminal law cannot make use of concepts of this sort to guide behavior. Nobody will be helpfully guided, for example, by a rule instructing them to commit offenses in *this* way rather than *that* way. Similarly, to instruct individuals that they must act in a certain way when they recognize their experience as being similar to some previous experience, a rule must characterize which previous experience they must employ, and the problem of describing that experience linguistically is simply the same problem that the appeal to recognitional concepts was designed to solve.

only about single actions; rather, they promulgate a universal rule that instructs courts to increase the sentence imposed on any offender whose offense satisfies the risk factor's definition. Because they apply to all offenses, the sentence enhancement imposed by any grading rule must reflect how that risk factor would affect the risk created by any possible crime in which it appears. Thus, considerably more information is required to promulgate a grading rule that grades by risk factor than simply to analyze the risk created by a single action: while evaluations of risk in a single instance need consider only the actual features of the actions under evaluation, lawmakers producing grading rules must identify how the risks created by quite different actions might all be affected by the presence of the same risk factor. Perhaps such investigations may be possible for some risk factors, especially relatively simple ones. But collecting that information will be far from trivial, and in many cases lawmakers will be unable to accurately quantify the risk factor's impact.

Furthermore, there is no reason to suppose that a particular risk factor will always produce an increase in risk of equal magnitude whenever it characterizes an action. That is, lawmakers might attempt to analyze the risk created by a particular risk factor by identifying the average increase it causes in the risk an offense creates, given some suitable method for weighting different kinds of offenses in computing that average. But this average risk increase would merely be a statistical construction that would not necessarily capture the increase in risk produced by a risk factor in any individual case. Instead, due to differences between the other features that characterize different offenses, the very same risk factor might differently influence the risk created by those different actions. Because causal relations are often extremely complex, the effect that any one feature of an action has on its likelihood of causing certain results often depends on other features of the same action: perhaps, for example, different methods of lighting a building on fire will be more or less dangerous depending on atmospheric conditions at the time of the fire, or perhaps different sorts of physical assault pose different risks of causing serious injury to victims of different gender or physical stature. The complexity of these relationships among different risk factors, in turn, poses a further challenge for defining grading rules: those rules can accurately quantify the impact each risk factor has on the risk

created by any particular offense only if they capture the causal interactions between different risk factors.

Grading by risk factor ordinarily treats the contributions to risk made by each risk factor as independent. Each grading rule, that is, gives courts two options: if an offense features a particular risk factor, the court must increase the sentence imposed by a specified amount, while otherwise it must leave an offender's sentence unchanged. Thus, each risk factor increases an offender's sentence by the same amount whenever it applies, regardless of any other features that characterize his offense. This uniform sentence increase is consistent with proportional punishment, however, only if risk increases uniformly, too: increasing the sentence of any offender by the same amount whenever a particular risk factor is present will estimate risk accurately only if the presence of that risk factor always increases the risk each offense creates by the same amount whenever it appears. But, as I have argued, the effect that one risk factor has on the risk an offense creates is unlikely to be independent of any other facts about that action; instead, different aspects of a single action will often interact causally, and thus the effect of one risk factor on the risk an action creates may differ depending on that action's other features. Grading rules that impose a uniform sentence enhancement whenever a particular risk factor appears, then, will simply grade offenses inaccurately. Perhaps they can correctly identify the average increase in risk attributable to a particular risk factor, but unless the effects of each risk factor are completely independent of one another, that risk factor will likely increase risk by a different amount (or may even decrease risk) in any individual case. The uniform sentence enhancement required by the grading rule, then, will not match the different increases in risk that the risk factor produces in the various offenses that it might characterize.

To accurately capture the actual effects of any given risk factor on the risk of individual actions, grading rules must not just reflect the effect that each risk factor has on all possible offenses it might characterize but must also capture the causal interactions between any combination of risk factors. In theory, such rules might be written: the law might produce one separate grading rule for each possible combination of risk factors, evaluate the increased risk produced by that specific



combination of risk factors, then individually specify the increased sentence to be imposed for performing an offense with that precise combination of risk factors. But in practice no lawmaker could actually produce such rules, which would in effect involve promulgating a separate grading rule for every different way in which an offense might be performed. Given how many different factors might affect the risk created by any given offense, the number of grading rules required would simply be enormous were a separate rule employed to grade each possible combination of risk factors. The task of promulgating that many rules would impose enormous burdens on lawmakers, who would be required to conduct separate empirical investigations of investigate the likelihood of harm arising from each possible combination of risk factors. It is extremely unlikely that lawmakers could actually develop a system of rules complex enough to reflect these causal relationships, but grading rules that failed to reflect them would inaccurately estimate the risk created by different offenses.

Thus, a system of grading by risk factor would likely be unable to define grading rules that estimate risk employing every risk factor that actually affects the risk offenses create. Even if it were possible in theory for such rules to be created, though, the resulting law of sentencing would be enormously complex, so much so that creating and administering it would be extremely impractical. Because grading rules apply prospectively, they must define risk factors in advance: under grading by risk factor, then, the law cannot merely evaluate the risk factors that characterize offenses that have already occurred but must rather promulgate a set of grading rules that identify and grade every risk factor that might in any way affect the risk created by any offense. A staggering number of rules would be required. Given sentencing law this complex, furthermore, adjudicating the sentences of individual offenders would likewise become enormously complex: since the sentence to be imposed would depend on which of these many risk factors characterized an offense, each court would be required to apply each of these innumerable grading rules to each offense. Sentencing offenders, then, would become a dramatically more difficult task. And in addition to posing practical challenges, that difficulty would itself be a source of inaccuracy: the more rules a court must apply, and the more complex those rules are,

the likelier it is simply to make a mistake. Of course, given adequate resources it might be possible to employ grading rules of such complexity, but the resources required would make grading by risk factor an unappealing alternative to other possible forms of grading. For example, the current Federal Sentencing Guidelines do include some rules that grade by risk factor, but because they include relatively few, they leave many risk factors unanalyzed.<sup>124</sup> And these guidelines are already frequently criticized for their length and complexity.<sup>125</sup> Grading rules that actually incorporated every risk factor would be vastly more complex. Even setting aside the challenges of identifying risk factors, then, it would be practically impossible for the law to actually employ a system of grading by risk factor complex enough to estimate risk accurately.

Thus, there are broad limitations on the scope of a system of grading by risk factor. Many risk factors simply cannot be employed in successful grading rules, either because lawmakers cannot promulgate a grading rule that verbally defines that risk factor and identifies a uniform increase it produces in the risk offenses create, or because courts would be unable to reliably identify when that risk factor characterized an offense. Furthermore, a system of grading rules that accurately captured all factors that affect the risk offenses create would be so complex as to be unworkable. Of course, these objections do not exclude all possible grading rules that grade by risk

---

<sup>124</sup> The sentencing guideline covering “Larceny, Embezzlement, and Other Forms of Theft; Offenses Involving Stolen Property; Property Damage or Destruction; Fraud and Deceit; Forgery; Offenses Involving Altered or Counterfeit Instruments Other than Counterfeit Bearer Obligations of the United States,” for example, is among the more complex in federal law, containing twenty different specific offense characteristics. U.S. SENTENCING GUIDELINES MANUAL § 2B1.1(b)(2) Among those are many that might be interpreted as forms of grading by risk factor, including whether the offense involved mass marketing, *id.* § 2B1.1(b)(2)(A)(ii), whether it involved theft from the person of another, *id.* § 2B1.1(b)(3), whether the offense involved the theft of medical products, *id.* § 2B1.1(b)(8), whether the offense involved misrepresentations in connection with charitable organizations, bankruptcy proceedings, or financial assistance for higher education, *id.* § 2B1.1(b)(9), whether the offense involved the knowing or intentional transfer of a trade secret to a foreign country, *id.* § 2B1.1(b)(14), whether the offense involved the possession of a dangerous firearm, *id.* § 2B1.1(b)(16)(B), whether the offense involved the dissemination of personal information, *id.* § 2B1.1(b)(18)(B), or whether the offense involved critical infrastructure, *id.* § 2B1.1(b)(19)(A)(i). Each of these risk factors, though, encompasses a wide range of conduct—there are many ways of employing mass marketing to defraud others, for example—and different forms of such conduct themselves likely differ in the degree of risk they create. And, of course, the risk factors explicitly enumerated in the guideline surely omit many other ways in which different offenses may differ in the risk they create of harm. Thus, even this relatively detailed grading rule is far from grading comprehensively by risk factor.

<sup>125</sup> *E.g.*, Frank O. Bowman, *The Failure of the Federal Sentencing Guidelines: A Structural Analysis*, 105 COLUM. L. REV. 1315, 1342–43 (2005) (noting the common criticism of the guidelines as too complex); Linda Greenhouse, *Guidelines on Sentencing Are Flawed, Justice Says*, N.Y. TIMES, November 21, 1998, <https://www.nytimes.com/1998/11/21/us/guidelines-on-sentencing-are-flawed-justice-says.html> (“Justice Breyer called the guidelines ‘simply too long and too complicated.’”).

factor, and in my view, the law should employ such grading rules when they are effective. But because there cannot be a grading rule for every feature that affects the risk an offense creates, actions can differ in the risk they create without differing in their risk factors, and thus residual differences in risk will persist between identically graded offenses. Consequently, grading by risk factor will fail, on its own, to grade offenses proportionally to risk: some offenders will be punished alike even though their offenses created different degrees of risk.

Because grading by risk factor will estimate the risk offenses create inaccurately, it will fail to be fully effective in guiding offenders' conduct. Of course, the grading rules the law does employ can effectively guide offenders to avoid the risk factors defined in those rules. But residual risk will remain: offenses that are identical with respect to the risk factors employed in grading may differ in the risk they create. And since no grading rule distinguishes between offenses that differ in this way, grading by risk factor does not guide offenders to reduce this sort of residual risk. If differences in residual risk were completely unknowable, it might be no objection that grading by risk factor cannot guide offenders to reduce them, since if they were entirely unknowable presumably no method of grading could successfully guide offenders to reduce them. But, in fact, offenders will often themselves be able to identify differences in the risk different offenses create, even though those differences are not attributable to any risk factor, for limitations on the effective scope of grading by risk factor depend on obstacles that confront courts and legislatures but do not confront offenders themselves.

First, courts can accurately apply only grading rules that employ risk factors that courts can reliably identify using their limited information about the manner in which each offense was performed. But because offenders have considerably better factual information than courts about how an offense was performed—after all, offenders were present during the crime itself and do not need to reconstruct it after the fact—they can identify factual differences between offenses about which courts are simply ignorant. Second, many difficulties that face grading by risk factor arise from the difficulty of defining universal, prospective rules for estimating risk: such rules must linguistically define the risk factor they employ, and they must specify the

degree to which that risk factor universally increases risk whenever it appears. By contrast, offenders themselves face a very different problem in estimating risk: they need only determine which of the possible actions they might perform are more or less risky. This latter task is considerably easier. In a wide range of contexts, individuals can often make correct decisions in guiding their own conduct even though they cannot verbally articulate any universal rule to govern their action: one can ride a bike successfully, for example, without being able to articulate verbal rules for successful bike riding, just as one can speak a language grammatically without being able to articulate the grammatical rules that govern it.<sup>126</sup> Offenders, many of whom may have considerable experience committing offenses, may thus be able to reliably compare the risk created by different methods of committing a particular crime—which physical actions are likelier to cause more severe injuries, say, or which interactions with crime victims are likelier to escalate into physical violence. If lawmakers cannot define any general rule that verbally defines a risk factor distinguishing two possible offenses and that quantifies the difference in risk between them, grading by risk factor cannot guide offenders in choosing between those offenses, even though offenders themselves nonetheless could recognize which is riskier. Offenders could reduce these kinds of risk, but grading by risk factor does not guide them to do so. While grading by risk factor can guide offenders to avoid the risk factors employed by whichever rules the law does employ, it cannot guide them to reduce residual risk not captured by any risk factor.

## 2. *Grading Directly by Risk*

Grading by risk factor, then, fails to account for the residual risk that distinguishes offenses to which the same grading rules apply. How else might the law estimate that residual risk in order to grade offenses accordingly? One obvious approach would be to grade directly by risk. Rather than promulgating separate grading rules that each correspond to a separate risk factor, this approach would simply promulgate a single grading rule stipulating that the sentence imposed for an offense should depend on its risk of producing a particular result. In sentencing each

---

<sup>126</sup> As Richard Heck puts it, “[T]here is a certain sort of gap between the reasons we have for our beliefs and the reasons we can communicate to others.” Heck, Jr., *supra* note 123, at 519.

offender, then, courts would directly estimate the degree to which his offense risked causing various harmful results—perhaps, say, by assigning each harm that an offense might have caused some probability between zero and one hundred. That risk estimate would then determine the sentence imposed: the likelihood of causing unintended harm would correspond directly to the magnitude of an enhancement in the sentence. On this approach, courts would perform each of the two tasks that grading by risk factor divides between courts and legislatures: they would first identify the facts about each offense to be employed in estimating its risk, then they would actually estimate it. Nonetheless, I will argue, because courts will likely err both in identifying the facts that ground risk estimates and in quantifying the risk each offense creates, grading directly by risk will likewise estimate risk inaccurately and thus guide conduct ineffectively.

In general, probability judgments are relevant to criminal activity only because offenders lack perfect information about the results of their offenses.<sup>127</sup> It is not certain that a particular offense will or will not cause any particular harm; rather, it is possible for any offense either to cause harm or not. Probability judgments attempt to quantify this uncertainty by measuring the fraction of possible states of affairs in which an offense actually does cause harm: if a particular offense has a one-fifth chance of causing death, say, then performing it actually will cause death in one fifth of all scenarios that are possible.<sup>128</sup> Employing these probability judgments to guide offenders' conduct will therefore reduce the total amount of harm produced. In any one case, of course, the actual state of affairs may be one in which harm does occur, but if offenders generally perform offenses that cause harms in fewer possibilities, then harm will occur in the actual state of affairs less often, and thus on the whole fewer harms will occur.<sup>129</sup>

---

<sup>127</sup> FRANK H. KNIGHT, *RISK, UNCERTAINTY, AND PROFIT* 218 (1921) (“The fundamental fact underlying probability reasoning is generally assumed to be our ignorance.”).

<sup>128</sup> *See, e.g., id.* at 212 (“[A]scertaining the numerical proportion of the cases in which X is associated with Y . . . yields the familiar probability judgment.”).

<sup>129</sup> Strictly speaking, this argument presupposes that each possibility is as likely to obtain as any other possibility. JOHN MAYNARD KEYNES, *A TREATISE ON PROBABILITY* 70 (1921) (“[A] numerical measure can actually be obtained in those cases only in which a reduction to a set of exclusive and exhaustive equiprobable alternatives is practicable.”); KNIGHT, *supra* note 127, at 220 (“[T]he alternatives themselves must be equally probable.”).

Obviously, which states of affairs count as possibilities will substantially influence any computation of the fraction of those possibilities in which harm occurs. Given complete information about the actual state of affairs, there would be only one possibility—namely, the single possibility compatible with that information.<sup>130</sup> And, of course, in that one possibility harm either occurs or does not.<sup>131</sup> Thus, probability judgments other than zero or one presuppose some limited set of information consistent with multiple different possibilities, and the probability of harm is the fraction among those possibilities in which harm occurs. What set of information should courts use to define the possibilities through which they estimate the risk an offense creates? One natural option would be to employ the information that the offender himself actually possessed. Offenders, after all, can deliberate using only the information they actually possess. Thus, if one possible offense is riskier than another only because of some fact of which an offender is unaware, it is not clear how a grading rule that estimates risk using that fact could guide conduct effectively. For without knowing that fact, the offender could not realize that the law would estimate that offense to be riskier and thus would punish it more severely. And offenders must recognize how different offenses will be punished in order for that difference to guide their conduct. Since offenders can guide their own conduct using only the information they actually possess, then, perhaps the law should estimate risk using only that information, so as to guide them to perform the action that is least risky given that information.

If sentences depend on the risk an action created, given the information that an offender actually possessed, though, the courts that sentence offenders must determine what information each offender actually did possess. And courts are unlikely to accurately reconstruct the actual beliefs each offender possessed when performing his offense. Of course, courts can reliably identify some such beliefs; indeed, to adjudicate questions of mens rea judges or juries must identify facts about an offender's mental state accurately enough to determine whether he acted

---

<sup>130</sup> KNIGHT, *supra* note 127, at 218 (“If it were possible to measure with absolute accuracy all the determining circumstances in the case it would seem that we should be able to predict the result in the individual instance . . .”).

<sup>131</sup> *Id.* at 219 (“And if the real probability reasoning is followed out to its conclusion, it seems that there is ‘really’ no probability at all, but certainty, if knowledge is complete.”).

knowingly, recklessly, or so forth. Mens rea ordinarily concerns mental states directed towards the intended results of actions, however, which usually can be inferred straightforwardly from a defendant's conduct.<sup>132</sup> If someone travels to a building, pours accelerant on it, and lights a match, no factfinder would struggle to determine he believed that he was performing those actions. But it is considerably more difficult to infer a defendant's beliefs concerning facts unrelated to his intentions, many of which will be crucial to determining the degree to which his action risked causing some unintended result. What evidence would show whether an arsonist recognized the materials of which a building was constructed, saw the dry grass surrounding it, or was paying attention to the weather forecast that morning? It is ordinarily relatively easy to infer an individual's intentions from his conduct precisely because intentions ordinarily manifest themselves in action, but because beliefs unrelated to intentions need not affect what an offender does, it is not clear how to infer them from his actions. Without accurately identifying an offender's actual beliefs, though, a factfinder cannot accurately estimate the risk an action created given the information an offender actually possessed.

The task of identifying factual information about each offense might be easier, of course, if the law estimated risk using a type of information that is easier to acquire than information about an offender's beliefs. In particular, the law might instead grade offenses based on the risk they create given information that was in some sense accessible to the offender, regardless of whether or not he actually possessed it. Certainly, information of this sort would often be easier to identify: factfinders would not need to investigate the offender's actual mental states and would instead need only to investigate the circumstances of the offense. But it is unclear how sentencing an offender according to a risk estimate that incorporates information he did not possess proportions the severity of punishment to the gravity of his wrongdoing rather than to his luck. Any fact that was unknown to an offender when he committed his offense might have been either true or false, and though some facts are likelier to be

---

<sup>132</sup> Some kinds of mens rea—particularly recklessness—may seem to require courts to identify a defendant's awareness of whether his action risked causing certain unintended results. In the second chapter of this dissertation, however, I will defend an account of recklessness according to which it does not involve a mental state of this sort, in part because I am skeptical courts can reliably identify a defendant's mental attitudes towards unintended results.

true than others, it seems purely a matter of chance which are true on any particular occasion. One offender, then, might receive a harsher sentence than another because of some factual difference between the two offenses that neither offender recognized. And if an offender lacked awareness of some particular fact, then that fact would affect neither the strength of the reasons he disregarded in acting nor the gravity of his wrongdoing. Allowing differences in facts unknown to offenders to affect their sentences, then, would allow differences in sentences not reflected in differences in the reasons different offenders disregarded in acting. And allowing facts of this sort to affect sentences appears to make the extent of criminal liability depend simply on luck. Although the task of estimating risk would be more feasible were courts not limited to the information offenders actually possessed, sentencing offenders according to that type of risk estimate seems inconsistent with the underlying motivations for grading by risk. If it is permissible for criminal sentences to depend on luck in this way, why not simply grade offenses by their results?

Furthermore, though courts would face an easier task in factually investigating each offense were they to estimate risk using the information accessible to offenders, not the information offenders actually possessed, they would likely perform even that easier task inaccurately. In general, the same sort of information would be required to estimate risk in this manner as would be required for courts to identify the risk factors that characterize an offense so as to impose a sentence under a scheme of grading by risk factor. And, much as I argued when analyzing grading by risk factor, courts will face challenges in performing this task accurately. No doubt, they can correctly identify some facts about each offense, particularly facts that can be identified through ordinary observation or through documentary evidence. But the risk an offense creates may depend on many details unlikely to be reflected accurately in evidence accessible to factfinders—for example, the risk that a violent crime will cause injury to persons or property may depend on precise details about how that violence is deployed, which witnesses may be able to describe only in broad strokes rather than in precise detail. Furthermore, the risk an offense created may depend on the circumstances in which it was performed. Evidence of those circumstances may sometimes exist, but oftentimes it will be unavailable: witnesses, for example, may



often be too distracted during criminal activity to notice details about their circumstances. And, of course, when only circumstantial evidence of the crime exists courts will struggle to identify any details whatsoever concerning the exact manner in which it was committed. Thus, though it may be easier to identify facts about an offense than to determine what beliefs an offender held, in neither case will courts always have access to all the information about an offense relevant to the risk it created. Instead, courts can employ only incomplete information, which will produce inaccurate estimates of risk.

The risk estimates courts make, then, will likely be based on incomplete or inaccurate factual information about each offense. Courts are unlikely to be more successful in actually producing those risk estimates. Under grading by risk, courts sentence offenders by quantifying each offense's likelihood of causing harm. Given the nature of probability judgments, it is difficult to see how courts ordinarily will be able to make reliable judgments about those likelihoods. As I have argued, probability judgments measure how often a particular fact is true in all possibilities that are compatible with certain specified information. Thus, to judge probabilities reliably requires some way of quantifying both all live possibilities and the fraction of such possibilities in which a particular result would occur. Roughly speaking, those possibilities might be quantified in one of two ways.<sup>133</sup> In some cases, a theoretical analysis of a particular action may be able both to enumerate all possible states of affairs and to identify those among them in which a particular result occurs. Games of chance—dice rolls, card draws, and the like—are susceptible to this sort of analysis: by understanding the physical symmetries of a die or the combinatorial structure of a card deck, we can identify the fraction of all possible die rolls or card draws that yield a particular result.<sup>134</sup> In other cases, by repeatedly performing a particular action we may be able to determine experimentally the fraction of possibilities in which that action produces a particular result.<sup>135</sup> This empirical approach characterizes the actuarial science used to price insurance products, say, or the quantitative analysis

---

<sup>133</sup> KNIGHT, *supra* note 127, at 214 (“There are two fundamentally different ways of arriving at the probability judgment of the form that a given numerical proportion of X’s are also Y’s”).

<sup>134</sup> *Id.* (“The first method is by a priori calculation, and is applicable to and used in games of chance.”).

<sup>135</sup> *Id.* (contrasting “the very different type of problem in which calculation is impossible and the result is reached by the empirical method of applying statistics to actual instances”).

underlying financial trading strategies.<sup>136</sup> It may be possible to estimate probabilities accurately when one of these two methods is employed to analyze a particular event.

Neither of these two approaches, however, will likely be available to courts estimating the risk created by individual offenses. Indeed, this limitation is not due to some peculiarity of criminal activity; rather, in general neither method of estimating probabilities may be applied to the common decisions of ordinary life. We lack knowledge of any theoretical principles that we could use to calculate a priori the fraction of total possibilities in which a particular offense causes a particular harm. Similarly, while statistical analysis might be applied to some broad categories of offenses—“burglary,” say, “armed robbery,” or “securities fraud”—the different offenses in that category would be enormously heterogenous, differing in respects that affect the likelihood that each would actually cause harm. To empirically estimate the risk created by one particular offense would instead require statistically analyzing some large set of instances identical to the actual offense—which does not and could not exist. As Frank Knight explains:

Take as an illustration any typical business decision. A manufacturer is considering the advisability of making a large commitment in increasing the capacity of his works. He “figures” more or less on the proposition, taking account as well as possible of the various factors more or less susceptible of measurement, but the final result is an “estimate” of the probable outcome of any proposed course of action. What is the “probability” of error (strictly, of any assigned degree of error) in the judgment? It is manifestly meaningless to speak of either calculating such a probability a priori or of determining it empirically by studying a large number of instances. The essential and outstanding fact is that the “instance” in question is so entirely unique that there are no others or not a sufficient number to make it possible to tabulate enough like it to form a basis for any inference of value about any real probability in the case we are interested in. The same obviously applies to the most of conduct and not to business decisions alone.<sup>137</sup>

Individual offenses are *sui generis*: they are sufficiently unlike other actions that techniques for estimating probabilities cannot fruitfully be applied to them. Of course,

---

<sup>136</sup> *Id.* at 215 (“It is difficult to think of a business ‘hazard’ with regard to which it is in any degree possible to calculate in advance the proportion of distribution among the different possible outcomes. This must be dealt with, if at all, by tabulating the results of experience . . . . It is evident that a great many hazards can be reduced to a fair degree of certainty by statistical grouping . . . .”).

<sup>137</sup> *Id.* at 226.

individual decision-making is not hopeless when confronted with uncertainty.<sup>138</sup> We do have some capacity to guide our conduct by foreseeing the results of our actions: in Knight's words, "We are so built that what seems to us reasonable is likely to be confirmed by experience, or we could not live in the world at all."<sup>139</sup> But this capacity to act reasonably in the face of uncertainty is not the capacity required in order to grade accurately by risk, which requires courts not to act but rather to precisely quantify the probability of harm.<sup>140</sup> Thus, quantitative estimates of the risk created by individual offenses will generally be inaccurate.<sup>141</sup> Grading offenses by those estimates of risk, then, would impose sentences that are disproportionate to the risk actually created by offenses.

This critique of grading by risk has focused on the difficulties courts would face in precisely quantifying the degree of risk created by an offense. If the inaccuracy of such estimates arises from the precision expected of courts in making them, though, perhaps grading by risk could be salvaged by allowing courts to estimate risk less precisely. Rather than assigning each harm potentially caused by an offense some likelihood between zero and one hundred, courts might instead merely sort offenses into two or three categories—say, based on whether they create a low, medium, or high risk.<sup>142</sup> Certainly, it is plausible that courts could more accurately classify offenses into these three categories than they could assign each a precise quantitative probability: since each category would include offenses differing widely in risk, an erroneous estimate of an offense's precise risk would often nonetheless fall

---

<sup>138</sup> *E.g.*, KEYNES, *supra* note 129, at 183 ("We may thus obtain information in cases where it would be impossible to ascribe any number to the probability in question.").

<sup>139</sup> KNIGHT, *supra* note 127, at 227.

<sup>140</sup> *Id.* at 230 ("The opinions upon which we act in everyday affairs and those which govern the decisions of responsible business managers for the most part have little similarity with conclusions reached by exhaustive analysis and accurate measurement.").

<sup>141</sup> KEYNES, *supra* note 129, at 182 ("It is evident that the cases in which exact numerical measurement is possible are a very limited class . . .").

<sup>142</sup> Some scholars have defended proposals along these lines. Alexander, Ferzan, and Morse, for example, propose that juries should classify the risk each offense creates into one of five categories, ranging from "virtually certain" to "virtually no risk." ALEXANDER, FERZAN, & MORSE, *supra* note 3, at 282. More generally, others have argued that offenders may be sentenced identically if their culpability is merely similar, though not identical. *See* Simons, *supra* note 23, at 1091–92 (defending the "principle of comparative culpability," on which marginal differences in culpability are not legally significant); Horder, *supra* note 28, at 769 (defending the "proximity principle," which holds that two offenders may be punished identically if there is "moral proximity" between their offenses). Perhaps this principle could be employed to justify sentencing offenders identically if their offenses differed only slightly in risk (though I am skeptical that risk classifications as course-grained as the one I have suggested here group together offenders whose culpability is similar enough to make the application of these principles plausible).

within the same category as the correct value, and thus the coarse-grained classification of that risk would still be accurate.

Nonetheless, although courts might sort offenses into relatively few categories of risk more accurately than they could precisely identify the exact risk each offense creates, grading offenses according to those coarse-grained classifications of risk would inaccurately grade offenses by their risk. Courts' coarse-grained classifications are more accurate than their precise risk estimates because coarse-grained classifications group together offenses that create very different degrees of risk: on a three-category classification of all offenses, an offense with a five percent chance of causing death and one with a thirty percent chance could both be classified as low risk. If offenses were graded according to this classification, those two offenses would be judged equally serious, and the offenders would be punished equally severely. But to grade these two offenses alike is simply to grade them inaccurately. A five percent chance of causing death is not, in fact, the same as a thirty percent chance of causing death; instead, the latter is six times greater than the former. And offense that creates a thirty percent chance of death is simply worse than one creating a five percent chance of death: in the aggregate, offenders who perform the former type of offense will produce six times as many deaths as those who perform the latter type, and it is surely six times worse for six times as many people to lose their lives. A method of estimating risk that classifies these offenses together just gets it wrong. Grading offenses according to a coarse-grained classification scheme does not solve the problem of inaccurate risk estimates; instead, it is an instance of the problem of inaccurate risk estimates, for it explicitly abandons the attempt to grade offenses accurately by risk due to the impracticality of doing so.

Grading rules that grade directly by risk, then, will grade offenses according to inaccurate risk estimates. Because courts' information about the manner in which offenses were performed will generally be limited, those estimates will not incorporate aspects of each offense that affect its likelihood of causing harm. Furthermore, courts are unlikely to estimate risk accurately even given that limited information: each individual offense is sufficiently unlike any other for precise risk estimates to be possible, and coarse-grained risk classifications do not even attempt to

grade accurately by risk. Grading by risk, like any grading rule, guides offenders' conduct through the sentences it imposes: by imposing a harsher sentence on offenses courts estimate to be riskier, the law guides offenders away from performing them. If courts estimate risk inaccurately, then, at least sometimes grading by risk will guide offenders to perform offenses that are not in fact less risky. In particular, each factor that undermines the accuracy of courts' risk estimates will similarly undermine grading by risk's effectiveness in guiding conduct.

First, because courts can use only information that they actually possess to estimate the risk that offenses create, grading by risk can guide offenders to reduce risk only in ways that will be reflected in the evidence that is available to courts. Nonetheless, courts will often lack evidence about many aspects of offenses that do affect their risk—precise details about exactly how a weapon was aimed, for example, or about exactly how physical force was applied to a victim's person or property. Because these aspects of offenses are unlikely to affect how courts estimate their risk and thus how courts sentence offenders, grading rules will not guide offenders to change these aspects of their offenses in ways that reduce the risk they create. Indeed, to some extent offenders themselves can control what evidence will exist for courts to review about the manner in which their offense was performed: offenders can destroy or create records of certain aspects of their offenses, for example, or they can influence which aspects of their conduct are or are not visible to witnesses. Thus, grading rules that allow the sentence imposed for an offense to depend so greatly on what information about the offense can be identified by the sentencing court may simply guide offenders to manipulate what evidence is available to the sentencing court rather than guiding them to reduce the risk their offenses create.

Second, courts' inability to accurately quantify the probabilities associated with each offense further undermines the effectiveness with which grading rules guide conduct. Because grading rules that grade by risk factor identify in advance how courts will estimate the risk that an offense creates, they allow offenders to deliberate about which potential offense to perform while knowing which ones would be sentenced more or less severely. By contrast, under grading by risk grading rules state only that courts must sentence offenders by estimating the risks their offenses

create; grading rules themselves do not identify which offenses will be judged riskier and thus punished more severely. Offenders, then, cannot simply consult legal rules to determine which of the offenses they might perform the law deems more serious and thus punishes more severely; instead, they must themselves estimate how courts will estimate the risk created by each possible offense. Since grading rules guide offenders through the punishments imposed for different offenses, grading by risk thus influences offenders to deliberate based on their own beliefs about how courts would estimate the risk different offenses create. Grading by risk can guide conduct effectively, then, only if offenders believe that courts will estimate the risk different offenses create correctly—that is, that a court’s quantitative estimate of the risk created by any offense will be greater than its estimate of the risk created by another just in case the former offense actually is riskier than the latter.

As I have argued, though, the risk created by individual events can be quantified accurately only when certain conditions are met—namely, when there exists some theoretical or empirical method of identifying the fraction of total possibilities in which harm would occur—and those conditions are not ordinarily met with respect to individual criminal offenses, or indeed with respect to any normal, everyday actions. The inability to quantify risk accurately does not undermine offenders’ own ability to guide their conduct: since offenders must choose between multiple possible actions, they need only compare those actions, and comparing the risk that different actions create is ordinarily considerably easier than quantifying that risk.<sup>143</sup> But a court sentencing an offender will consider only the actual offense that he performed: it must explicitly quantify the risk that offense created rather than comparing it to some other offense that the offender might have performed instead. Thus, whether the sentences offenders expect courts to impose for different offenses correspond to differences in the actual risk those offenses create depends not on the accuracy of some comparative judgment—as, under grading by risk factor,

---

<sup>143</sup> KEYNES, *supra* note 129, at 182–83 (“In actual reasoning . . . exact numerical measures, will occur comparatively seldom. The sphere of inexact numerical comparison is not, however, quite so limited. . . . This method is frequently adopted in common discourse. When we ask how probable something is, we often put our question in the form—Is it more or less probable than so and so?—where ‘so and so’ is some comparable and better known probability. We may thus obtain information in cases where it would be impossible to ascribe any number to the probability in question.”).

legislatures need only compare the risks created by offenses performed with or without a particular risk factor—but rather on the accuracy of courts’ quantitative estimates of risk. And in the absence of any basis with which to compute the fraction of possibilities in which harm would occur, those estimates will be hopelessly inaccurate. Of course, offenders may not be able to predict how courts’ estimates will err; instead, they should expect courts’ errors to be largely random. But if offenders expect that the sentence imposed for each possible offense will be randomly and unforeseeably disproportional to the risk it creates, grading by risk will simply provide no guidance whatsoever, for offenders will generally have little reason to expect offenses that actually are riskier to be graded correctly as being riskier and to be punished accordingly.

To be sure, perhaps sometimes offenders should expect differences between courts’ estimates to correspond to differences between the actual risk offenses create. In particular, if offenses differ wildly in the risk they create—an unarmed robbery, say, versus a robbery begun by firing multiple warning shots into a crowd—it might be reasonable to expect that a court’s estimate would be higher for the risk created by the riskier offense, even if each estimate on its own would still be quite inaccurate. Similarly, if grading by risk employed course-grained risk classifications, offenders might reasonably expect courts often to accurately classify offenses into those categories, and thus might expect greater sentences to be imposed for high-risk offenses than for low-risk ones. Thus, grading by risk might still guide offenders to prefer offenses that create dramatically less risk than the alternatives do. But many possible offenses differ in the risk they create in less substantial ways. Of course, the importance of reducing risk is greater for greater degrees of risk, but reducing risk even by smaller quantities still reduces the harm that criminal offenses will in the aggregate produce. Grading rules ought to guide offenders to reduce risk to any extent possible, not only when risk can be reduced at one stroke by a substantial amount. Because it grades offenses using inaccurate risk estimates, however, grading by risk cannot effectively provide that guidance. As with grading by risk factor, many differences that affect the risk offenses create will not reliably be reflected in

differences in the sentences imposed for performing them, and thus grading by risk will not effectively guide offenders to choose offenses that in fact create less risk.

### *C. Risk Estimates by Offenders*

My critique of grading by risk has focused on how grading rules discharge a particular function—namely, guiding offenders in deliberating over how to perform criminal acts. That deliberation takes place from offenders’ own perspectives. Each offender possesses certain information about his circumstances and about the various offenses he might perform; then, using that information, he must choose among those offenses. Risk estimates constitute an intermediate step in this deliberative process, one that depends upon the information with which the offender began deliberating, and upon which the decision that concludes that deliberation in turn depends. Since it is in society’s interest for offenses generally to cause as few unintended harms as possible, the law ought to guide each offender to perform the offenses that are least likely to cause harm. Thus, each offender should choose how to offend at least in part by considering the risk each offense would create. Those risk estimates must be made, furthermore, using the information offenders actually possess. If offenders accept that they must choose the option that is least risky according to some method of computing risk, they are committed to always choosing the same option in any choice made based on the same information, for absent any difference in the information on which deliberation is based, nothing could justify a different choice.<sup>144</sup> Offenders should choose the method of computing risk that produces the least harm when employed generally. But since the choice of a particular option commits offenders to make the same decision in any context in which they possess the same information, the choice that will produce the least harm in the aggregate is the choice that causes harm in the fewest scenarios consistent with the information that an offender actually possesses. And, of course, that option is just the option whose risk is lowest when measured relative to an offender’s actual information.

Grading by risk judges the gravity of an offense’s wrongdoing by the risk it creates of unintended harm. The notion of risk relevant to those judgments is firmly

---

<sup>144</sup> See *supra* Part I.A.1.a.



anchored in the offender's own perspective both in its origin and in its use: the risk each offense would create must be evaluated based on the information the offender actually possesses, and those estimates must in turn guide the offender's own choice of which offense to perform. But although estimating the risk each offense creates forms an intermediate stage in the offender's own deliberations, under grading by risk the authoritative estimates of risk are produced not by the offender himself but rather by the government entities that estimate risk in order to impose sentences. Grading by risk distributes the process of deliberation and choice between the offender and those government entities. Estimating risk is not simply an intermediate step that the offender himself performs as his deliberation proceeds from facts about each option through assessments of the risk each creates to a choice of one to perform. Rather, the law, not the offender, estimates the risk each offense creates in order to grade it; then, since those risk estimates determine how each offense would be sentenced, they guide offenders in selecting among possible offenses. Though deliberation both begins and ends with offenders, under grading by risk a central step in that deliberation is performed by the state.

By distributing this process among multiple different entities, grading by risk introduces a new kind of problem—communication. If different steps in deliberation are performed by different entities, then somehow the results of each step must be communicated from the entity that performed that step, and thus produced those results, to the different entity that must perform the next step in deliberation based on those results. If courts must estimate risk using offenders' own information, then that information must somehow be communicated from offenders to courts to be employed in estimating risk. Similarly, for grading rules to guide offenders to perform less risky offenses, the risk estimates that the law employs to sentence offenders must somehow be communicated from the government entities that create them to the offenders who must employ them. In my critique of grading by risk, I identified two general sources of inaccuracy in the risk estimates it produces, each of which undermines the effectiveness with which it guides conduct. Those two sources of inaccuracy correspond to the two types of communication made necessary because grading by risk distributes deliberation between offenders and government bodies.

Communicating each type of information is difficult, and the methods through which grading by risk may do so meet with only limited success. These difficulties in communication explain why grading by risk guides conduct ineffectively.

Under both grading by risk factor and grading directly by risk, courts must determine the manner in which each offense was performed in order to estimate the risk that it created. This task involves communication between offenders and courts: offenders themselves possess the information that must be employed to determine how risky each offense is, and that information must be communicated to the courts that actually estimate risk. As I have argued, though, courts will generally perform this task inaccurately. It is difficult to retrospectively reconstruct the exact manner in which an offense was performed, which involves a large number of precise details; in many cases no evidence of such details will exist at all, and in other cases that evidence will fail to distinguish precisely between similar possible offenses. Thus, courts will likely fail to identify many details relevant to the risk that offense created. Grading by risk requires information to be transmitted from offenders to courts; it will estimate risk inaccurately, and thus guide conduct ineffectively, because that information often will be transmitted incompletely or inaccurately.

Since offenders themselves must ultimately choose which offense to perform, the law's estimates of the risk various offenses create must then be communicated back to offenders: somehow the law must inform offenders which of the various offenses they might perform would create more or less risk, so that offenders may choose the less risky ones. Grading by risk factor and grading directly by risk differ in the strategies they take to identify more or less risky offenses. Nonetheless, neither effectively informs offenders of which of their options create relatively more or less risk. Once again, then, grading by risk makes necessary a particular kind of communication because it requires offenders and various government entities to jointly carry out the process of deliberating over which offense, but it cannot effectively transmit the information required for that distributed deliberation to be effective.

Grading by risk factor attempts to decompose the risk each offense creates into separate components each contributed independently by the different risk factors

that characterize each offense. For each such risk factor, the legislature promulgates a separate grading rule that specifies the increased punishment to be imposed on any offender whose offense features that risk factor. Offenders, in turn, can use those grading rules to identify which offenses the law judges to be riskier. The law expresses how it judges risk through the sentence enhancements grading rules require, which offenders may calculate by identifying which risk factors characterize each offense, then summing the increased sentences those risk factors would produce. If grading by risk factor did identify a set of risk factors that each increased the risk created by any offense by some specified amount, and if the increased sentences required by each grading rule did accurately quantify the increase in risk attributable to each risk factor, then this approach would accurately inform offenders of how to compare the risk their options created. That is, if the increased sentence due to each risk factor corresponded to the difference in risk between any two offenses that differed only with respect to that risk factor, then one offense would be riskier than another just in case grading it by risk factor sentenced it more severely. And offenders could use the sentences imposed by those rules to identify more or less risky offenses. But practically speaking lawmakers cannot identify a set of independent risk factors or accurately quantify the increase in risk attributable to each. Furthermore, any set of rules that did grade according to all such risk factors would be so complex as to be unusable by actual offenders. Thus, grading by risk factor does not successfully communicate to offenders which of their options creates the least risk, and does not guide their conduct effectively.

Grading by risk factor attempts to inform offenders of which offenses are riskier through grading rules that identify in advance the features that make offenses riskier. These rules directly communicate to offenders how they should choose among their options. By contrast, grading directly by risk only indirectly informs offenders which options they should choose. The risk estimates that courts produce determine how severely each offense will be sentenced, thereby identifying the seriousness of different offenses. But since courts do not estimate the risk an offense creates until after it is performed, during their own deliberations offenders obviously cannot rely on courts' actual estimates to determine which available option would be least

serious. Instead, offenders must deliberate and choose based on their beliefs as to how a court would evaluate the risk created by each offense that they might perform. Thus, the method of estimating risk that courts employ guides offenders' conduct by influencing their beliefs about which of the offenses they might perform would be punished more severely.

In deliberating, offenders face a choice between multiple options; to guide that choice they require some principle for comparing the risk those options create. Grading directly by risk, however, estimates the risk created by only one offense at a time—namely, whichever offense actually is chosen. These individual estimates do not, in themselves, communicate any comparative information about the relationship between different options. Thus, courts must employ a method of separately estimating the risk each offense creates that enables offenders to compare the expected risk estimates for each possible offense. The method by which to do so is obvious. In general, we compare different things by establishing some common standard against which they may each be individually measured. For example, we compare the lengths of different objects that cannot be compared directly by measuring each against some common standard—the meter stick, say—whose length is constant. If the risk each offense created could be measured by comparing it accurately with some common standard kept constant between each measurement, then the risk estimates courts produce would be greater for offenses that are actually riskier, thereby guiding offenders' conduct correctly.

Unfortunately, however, no usable common measure exists to measure risk precisely. The most natural suggestion—quantitative probability estimates—uses certainty as a common standard: such estimates measure a risk of harm by precisely measuring how much less likely that harm is than a harm that is certain to occur.<sup>145</sup> If courts were capable of identifying exactly how much less than certain a particular action was to cause harm, this approach would guide conduct effectively, since only if one offense is less risky than another is the former less than certain by more than the latter to cause harm. But outside of unusual conditions rarely met in criminal

---

<sup>145</sup> See, e.g., KEYNES, *supra* note 129, at 182 (precisely defining numerical probability as a particular relation between the probability of a proposition and certainty).

courtrooms, precise comparisons of likelihoods with certainty are extremely unreliable. Thus, courts' judgments of how much less than certain are the risks created by offenses are unlikely to will correctly inform offenders which offenses are riskier.

Alternative common standards exist, of course. Rather than using precise quantitative metrics, each action might be evaluated according to some intuitionistic standard that roughly classifies actions into relatively few groups according to their risk. These intuitionistic judgments are likely to be more accurate than attempts at precise quantification, but because they are intuitionistic, they will be vague, classifying actions as alike in risk even when the risk they create differs. Thus, offenders cannot expect that courts employing intuitionistic standards to measure risk will estimate one action to be riskier than another whenever it is, in fact, riskier. Because grading directly by risk guides offenders to act based on how they believe courts will estimate risk, it cannot effectively identify the offenses offenders should avoid unless offenders expect courts' risk judgments to accurately reflect differences in the risk offenses create. But because courts lack a workable common standard with which to measure that risk, offenders cannot rationally expect differences in its judgments of the risk offenses create to accurately track differences in the risk they actually do create. Like grading by risk factor, grading directly by risk fails to correctly inform offenders of which offenses they should perform.

Thus, the ineffectiveness of different forms of grading by risk has a common source. Estimates of the risk different offenses would create play a key role in an offender's deliberation: they depend on his own information about his possible actions, and they guide his choice among those actions. Because grading by risk requires government entities to actually produce those risk estimates, though, it must solve two difficult communication problems—namely, communicating information about each offense from offenders to courts, then communicating judgments about the relative risk created by different actions from some government entity back to offenders. Neither form of grading by risk solves these problems; thus, neither guides conduct effectively. If that ineffectiveness arises from the fact that grading by risk distributes deliberation among multiple entities so that communication between them

becomes necessary, the sort of change that would avoid these problems might seem obvious: offenders' deliberation should rely not on some government entity's risk estimates but rather on their own risk estimates.

Grading by risk makes communication necessary because the government entity that estimates risk differs from the offender whose information must determine those risk estimates and who must employ them in choice. But no communication is needed for an offender to estimate risk using information that he himself possesses, nor for him to use those risk estimates to choose among possible offenses. And if no communication need occur at all, inaccuracies in communication cannot undermine the effectiveness with which the law guides offenders to choose less risky offenses. Offenders themselves will automatically use the information they actually possess to estimate risk, whereas errors arise if information must be transmitted from offenders to some other entity that estimates risk. Similarly, offenders themselves can directly compare the options between which they must deliberate and choose, whereas all methods grading rules might use to communicate comparative information about risk created by different offenses introduce inaccuracies into those comparisons. The problems with grading by risk may be avoided if the law does not itself estimate risk but rather requires offenders to do so.

On this approach, criminal law would promulgate rules that instruct offenders to choose which offense to perform based on their own beliefs about those offenses. At first glance, regulating individuals' conduct according to their own beliefs might seem odd. But in fact this is a common method for criminal prohibitions to employ in guiding conduct. Indeed, perhaps the most central element of criminal law also guides conduct through the very same approach—namely, the criminalization rules that define offenses using intentional forms of mens rea. While grading rules aim to reduce harm by guiding offenders to perform less dangerous offenses, criminalization rules aim to entirely prohibit certain actions that are excessively likely to cause harm. Just as grading rules must identify which offenses are more or less risky, then, criminalization rules must identify which actions are likely enough to cause harm that they should not be performed. But criminalization rules do not themselves identify such actions; rather—at least when defining intentional crimes—they instruct

individuals not to perform actions that individuals themselves judge sufficiently likely to cause harm, just as, on my suggestion here, grading rules could instruct offenders to reduce the risk their offenses create according to their own judgment of those risks.

In general, since we care about the results of our actions, we decide what to do in part based on what we expect would happen if we chose to perform different actions. Thus, ordinarily an important part of deliberation consists of identifying at least some of the results that would be produced by each action we might choose. This element of deliberation does not occur because the law demands it; rather, we identify the results of our potential actions because those results are often relevant to our own purposes. But the law can exploit this element of deliberation for its own purposes. That is, if in deliberation we ordinarily inquire into at least some of the results of our actions, the law can regulate our conduct based on the results of that inquiry: in particular, it can prohibit individuals from performing actions that they themselves judge will cause certain harmful results. If individuals have already identified which actions are likely to be harmful, the law need not investigate that question itself; instead, it can simply instruct individuals to avoid actions based on the conclusions they themselves have already drawn about those actions' results. Indeed, for all the reasons I have canvassed in criticizing grading by risk, individuals' own judgments as to whether a particular action will cause harm are likelier to be more accurate than any method the law might employ on its own to evaluate and specify which actions are so risky that they must be avoided entirely.<sup>146</sup> Thus, prohibiting individuals from performing actions that they conclude in their own planning will cause harm will generally reduce harm more effectively than prohibitions relying on the law's own specification of which actions are harmful. And, broadly speaking, the results of an action that form part of the agent's plan in acting are those that he causes intentionally.<sup>147</sup> By prohibiting individuals from intentionally causing certain harms,

---

<sup>146</sup> Of course, sometimes the law can identify dangerous actions more accurately than individuals can, and it can directly prohibit those actions—traffic offenses, gun offenses, drug offenses, and the like—rather than simply defining offenses as actions that cause certain harms. See, e.g., R.A. Duff, *Criminalizing Endangerment*, in *DEFINING CRIMES: ESSAYS ON THE SPECIAL PART OF THE CRIMINAL LAW* 43, 59 (R.A. Duff & Stuart Green eds., 2005); A.P. Simester & Andrew von Hirsch, *Remote Harms and Non-constitutive Crimes*, 28 *CRIMINAL JUSTICE ETHICS* 89–107, 95 (2009); Peter Westen, *The Ontological Problem of "Risk" and "Endangerment" in Criminal Law*, in *PHILOSOPHICAL FOUNDATIONS OF CRIMINAL LAW*, 308 (R.A. Duff & Stuart P. Green eds., 2011).

<sup>147</sup> E.g., J.L. Austin, *Three Ways of Spilling Ink*, 75 *PHIL. REV.* 427, 437 (1966) ("I must be supposed to have as it were a plan, an operation—order or something of the kind on which I'm acting, which I am seeking to put into

then, criminal law requires individuals to employ their own assessments of the likelihood that different actions will cause harm in deliberating over whether to perform those actions.

By defining prohibited conduct using intentional grades of mens rea, criminalization rules guide offenders to exclude certain actions from deliberation according to their own beliefs about whether those actions would result in harm. Perhaps, then, by grading offenses using mental states grading rules could similarly guide offenders to choose among different offenses based on their own estimates of the risk those offenses would create. Nonetheless, I will argue that grading by mental state will not be effective. The alternative I propose should be unsurprising: grading by result. Under grading by result, of course, grading rules entirely abandon any attempt to proportion the sentence imposed on each offense to the risk it individually creates, however that risk might be measured. Grading by risk in some form certainly seems like the most natural way to guide offenders to perform less risky offenses, but because the law cannot estimate risk accurately grading by risk cannot, in fact, effectively guide conduct. By contrast, the results of actions are considerably easier to accurately identify: courts will not generally struggle to determine whether a particular crime produced a particular harmful result. And, I will argue, unlike forms of grading by mental state, which unsuccessfully attempt to directly instruct offenders to choose offenses based on their own estimates of the risk those offenses create, grading by result succeeds at indirectly guiding offenders to do so. Thus, I will conclude, grading by results achieves what grading by risk cannot—it effectively guides offenders to reduce the risk their offenses create.

### *1. Grading by Mental State*

Which mental states could grading rules employ in order to grade offenses? Perhaps the most obvious suggestion would be the various mental states required by different grades of mens rea, mental states that the law already deems relevant to a defendant's culpability. Intentional forms of mens rea, of course, could not be

---

effect, carry out in action: only of course nothing necessarily or, usually, even faintly, so full blooded as a plan proper. When we draw attention to this aspect of action, we use the words connected with intention.”); *see also* DUFF, *supra* note 9, at 56 (“But I must believe that there is some chance that my action will have the desired result if I am to count as intending that result.”).



employed by grading rules concerned with the risk of unintended harms. But other kinds of mens rea—recklessness and negligence—are concerned with unintended harms, and thus grading rules could evaluate offenses based on the defendant’s recklessness or negligence. Indeed, some proposals exist to grade offenses by only those harms towards which offenders are reckless or negligent.<sup>148</sup> On this approach, then, grading rules would grade by mens rea. But although kinds of mens rea like recklessness and negligence clearly are relevant in some way to culpability, I will argue that because criminalization rules and grading rules serve different functions, the mental states required by those kinds of mens rea cannot simply be repurposed for the use of grading rules. Consequently, grading rules that grade by negligence or by recklessness will not guide conduct effectively.

Alternatively, grading rules might employ a type of mental state *rea* distinct from the mens rea employed by criminalization rules. In particular, if grading rules are concerned to identify riskier offenses as more serious, the most relevant sort of mental state would seem to be the offender’s own estimate of the degree of risk that his offense created. Such grading rules would grade by risk, but they would not require a state entity to produce risk estimates according to some legal procedure. Rather, they would simply defer to offenders’ own judgments, grading offenses by the offender’s own risk estimates. But although the mental states employed by this proposal are at least suited to the distinctive function of grading rules, I will argue that they too will be ineffective. First, actually attempting to employ rules of this sort to evaluate particular offenses would be unworkable, practically speaking. To grade by offenders’ risk estimates, a court must be able to identify the particular degree of risk that an offender believed his offense would create, but I will argue that in general courts cannot reliably identify beliefs of this sort. Second, I will argue that grading rules that grade by offenders’ risk estimates cannot effectively guide them to avoid causing unintended harms even in theory. Individuals typically do form some estimate of the likelihood that acting will produce an intended result, since those estimates are relevant to whether an action will achieve the agent’s own purposes in

---

<sup>148</sup> *E.g.*, MODEL PENAL CODE § 210.1(b) (AM. LAW INST. 1962) (proposing the replacement of felony murder with a rebuttable presumption of recklessness if death is caused during certain enumerated felonies).

performing it. But since unintended results are, by definition, independent of an offender's purposes in acting, in general he need not estimate the likelihood that they will occur. And if individuals do not make those risk estimates, grading rules that grade by those risk estimates cannot distinguish among the seriousness of different offenses and thus cannot guide individuals' conduct at all. To reduce unintended harms, the law must not merely instruct individuals to choose less risky offenses, once the risk those offenses would create has been estimated, but must further guide individuals to actually make those estimates in the first place. Because grading by offenders' own risk estimates provides only the latter instruction, then, not the former, it will not effectively guide offenders to reduce risk.

*a. Grading by Mens Rea*

Grading by mens rea has often been proposed as a response to the objections raised against grading by result. English lawyers have extensively debated what has been called the correspondence principle, a version of grading by mens rea that would require some type of mens rea to correspond to every element defining the actus reus of a crime.<sup>149</sup> And a principle of this sort has formed the basis of many actual proposals to reform criminal law. The Model Penal Code, for example, replaced the traditional felony murder rule with a rebuttable presumption that committing certain felonies evinces the reckless indifference required for murder,<sup>150</sup> although many jurisdictions that generally adopted the Code have rejected this particular provision.<sup>151</sup> Scholars generally skeptical of the criminal law's use of results have seen this proposal as a promising avenue for reform.<sup>152</sup> Similarly, in compliance with the correspondence principle a number of proposed reforms to English law have eliminated offense definitions on which a difference in results alone, with no

---

<sup>149</sup> For a definition of the principle, see Ashworth & Campbell, *supra* note 35, at 192. For discussion, see Ashworth, *supra* note 6; Ashworth & Campbell, *supra* note 35; Gardner, *supra* note 28; Horder, *supra* note 28; Jeremy Horder, *Questioning the Correspondence Principle—A Reply*, 1999 CRIM. L. REV. 206; Adam Jackson & Tony Storey, *Reforming Offences Against the Person: In Defence of 'Moderate' Constructivism*, 79 J. CRIM. L. 437 (2015); Barry Mitchell, *In Defence of a Principle of Correspondence*, 1999 CRIM. L. REV. 195.

<sup>150</sup> MODEL PENAL CODE § 210.2(1)(a) (requiring the mens rea necessary for a murder conviction to be “presumed if the actor is engaged or is an accomplice in the commission of, or an attempt to commit, or flight after committing or attempting to commit robbery, rape or deviate sexual intercourse by force or threat of force, arson, burglary, kidnapping or felonious escape”).

<sup>151</sup> Guyora Binder, *Felony Murder and Mens Rea Default Rules: A Study in Statutory Interpretation*, 4 BUFF. CRIM. L. REV. 399, 400–01 (2000).

<sup>152</sup> Schulhofer, *supra* note 3, at 1595–97.

corresponding difference in mens rea, can increase the seriousness of an offense,<sup>153</sup> although these reforms have not been adopted.<sup>154</sup> These proposals, then, would grade offenses by their result, but in a modified form: rather than grading by results alone, grading rules, like criminalization rules, would allow only those results towards which a defendant possessed some form of mens rea to affect the scope of his liability.<sup>155</sup>

Nonetheless, implementing grading by mental state through the correspondence principle simply fails to address the central objection to grading by results—namely, that it imposes disproportional punishments on offenders. Its critics argue that grading by results punishes disproportionately by requiring otherwise identical offenders to be punished differently if their offenses produce different results. But merely adding a mens rea requirement to grading rules does not avoid this problem: in effect, the mens rea requirement ensures that the law grades by results only if a defendant possessed mens rea towards that result, which is obviously still a form of grading by results. To be sure, if some defendants lack mens rea towards some results of their offenses, those results would not affect sentences under grading by mens rea: all offenders who lack that additional mens rea would be punished alike regardless of whether their offenses cause a harmful result. Thus, perhaps grading by mens rea grades by fewer results than grading by results would. But under grading by mens rea all offenders who do possess the required form of mens rea are graded by result: an offender who causes a harmful result with mens rea would be punished more severely than one with the same mens rea who fails to cause harm. Thus, a mens rea requirement prevents the law from sentencing offenders by results only if they

---

<sup>153</sup> See, e.g., LAW COMMISSION, CONSULTATION PAPER NO 217, REFORM OF OFFENCES AGAINST THE PERSON: A SCOPING CONSULTATION PAPER 71 (2014) (criticizing existing law on the grounds that “in many offences, there is a lack of correspondence between the external elements of the offence and the required mental element. That is to say, D can be held liable for results that are more severe than what D intended or foresaw.”); Law Commission, *Home Office 1998 Draft Bill*, in LAW COMMISSION NO 361, REFORM OF OFFENCES AGAINST THE PERSON 212, 1–4 (defining four offenses of, essentially, intentionally or recklessly causing injury or serious injury to another). For criticism of these proposals, see generally Gardner, *supra* note 28.

<sup>154</sup> *Offences Against the Person: Current Project Status*, LAW COMMISSION (Mar. 16, 2022), <https://www.lawcom.gov.uk/project/offences-against-the-person/>.

<sup>155</sup> On its face the correspondence principle applies only to the definitions of criminal offenses. Nonetheless, it could ground a more general approach to grading on which offenders must possess mens rea towards any result that affects the scope of their liability, regardless of whether it affects their liability by changing the offense of which they were convicted or by affecting some evaluation of the seriousness of that offense, which in turn affects its sentence.

lack mens rea; it leaves the sentencing of offenders who possess mens rea unchanged. And if grading by results is objectionable, why would it become acceptable when applied only to offenders who possess mens rea? If results alone cannot justify punishing otherwise identical offenders differently, why can they justify punishing identical offenders with mens rea differently?

Although the correspondence principle limits which results the criminal law may employ in grading, even were it accepted grading by mens rea would still continue to grade offenses by their results. Thus, though the correspondence principle is often defended as an alternative to grading by results, it in fact runs afoul of one primary objection raised against grading by results. Nonetheless, grading by mens rea might easily be modified to avoid this objection: it could simply stop grading offenses by results entirely. On this approach, the law would sentence defendants based on only whether they possess mens rea with respect to certain harms. Indeed, the law already grades some offenses in this manner: the Model Penal Code, for example, grades interference with the custody of children as a third-degree felony, rather than a misdemeanor, if the defendant “acted with knowledge that his conduct would cause serious alarm for the child’s safety, or in reckless disregard of a likelihood of causing such alarm.”<sup>156</sup> Kinds of mens rea like purpose and knowledge are directed towards the intentional results of actions; the grades of mens rea that offenders might possess towards the harms their actions unintentionally risk are those grades of mens rea that define unintentional crimes, such as recklessness and negligence. Perhaps grading by mental state could guide offenders to reduce the risk of certain harms by sentencing offenders more severely if they are reckless or negligent with respect to those harms.

But because criminalization rules and grading rules have different functions, the kinds of mens rea used in criminalization rules cannot simply be directly transferred to grading rules. Indeed, the inapplicability of kinds of mens rea like recklessness or negligence to grading rules can be read off the definitions of those kinds of mens rea themselves. The Model Penal Code, for example, provides that recklessness involves the disregard of a risk “of such a nature and degree that . . . its disregard involves a gross deviation from the standard of conduct that a law-abiding

---

<sup>156</sup> MODEL PENAL CODE § 212.4(1) (AM. LAW INST. 1962).

person would observe in the actor's situation."<sup>157</sup> Negligence, similarly, involves the failure to perceive a risk "of such a nature and degree that the actor's failure to perceive it . . . involves a gross deviation from the standard of care that a reasonable person would observe in the actor's situation."<sup>158</sup> Grading by mens rea would grade individual offenses by asking whether offenders possessed these types of mens rea with respect to certain harms—for example, if mens rea were required for felony murder, conviction would require a defendant to have been reckless or negligent towards the risk of causing death. The jury trying, say, an arsonist whose arson results in accidental death might be asked to determine whether the manner in which the defendant burned down a building deviated from the standard of conduct of a law-abiding person; another jury might be asked whether a rapist who accidentally kills his victim employed a reasonable standard of care in perceiving the risks posed by his rape.

How, exactly, does the law-abiding arsonist burn down a building? What, exactly, is a reasonable standard of care to employ during rape? Even to ask these questions is to perceive their absurdity. The law-abiding arsonist and the reasonable rapist are contradictions in terms; rape and arson are serious felonies prohibited by law, and, for precisely that reason, reasonable, law-abiding people do not commit them. Grading by risk grades offenses based on whether an offender possessed an additional type of mens rea, in addition to whichever mens rea satisfies the underlying offense. But the definition of that additional required form of mens rea will be satisfied automatically in virtue of the defendant's having committed the underlying offense, since the forms of mens rea at issue—recklessness and negligent—are on their face satisfied by the intentional commission of a serious felony. Merely by committing arson or rape one violates the standard of conduct of law-abiding persons or the standard of care of reasonable people; there is no way to commit those crimes without violating those standards.

---

<sup>157</sup> *Id.* § 2.02(2)(c).

<sup>158</sup> *Id.* § 2.02(2)(d).

The problem, furthermore, is not merely a practical issue of inflexible statutory language requiring technical redrafting.<sup>159</sup> Instead, ordinary concepts of recklessness and negligence cannot sensibly be applied to unintended results caused during other crimes because those concepts play a different role within criminal law: the types of mens rea employed by criminalization rules define the boundary between prohibited and permissible conduct, which is simply a different task than comparing the seriousness of different forms of prohibited conduct. In particular, forms of mens rea like recklessness and negligence identify which risks individuals are prohibited from taking. Virtually every activity performed in everyday life creates some degree of risk, but precisely because those risks are inevitable, the law cannot hope to eliminate them entirely. Rather, at best the law can only reduce the risks we impose on one another to an acceptable level. Some risks are sufficiently small to be outweighed by the benefits of the actions that create them, and those actions are permitted, lest we sacrifice the benefits they produce. As risk increases, though it eventually becomes sufficiently great to outweigh those benefits; at that point, we must refrain from acting. Recklessness and negligence mark the point at which a risk becomes large enough that we are prohibited from creating it. Criminalization rules specify which actions are prohibited and which are permitted; offenses defined in terms of recklessness or negligence distinguish the two based on the magnitude of the risk. Indeed, the Model Penal Code defines each form of mens rea in these terms: the risk that a reckless defendant disregards and that a negligent defendant fails to perceive must be “substantial and unjustified.”<sup>160</sup> Creating such risks—those too great for any reasonable or law-abiding person to create—are forbidden.

But this question—when an action is risky enough to be prohibited—is completely irrelevant to the question that grading rules must answer. Recklessness and negligence address questions of criminalization: they distinguish permissible forms of risk-taking from those that are prohibited because the risk created was

---

<sup>159</sup> Cf. Kevin Cole, *Killings during Crime: Toward a Discriminating Theory of Strict Liability*, 28 AM. CRIM. L. REV. 73, 108 (1990) (expressing concern that, given the Model Penal Code’s definitions of recklessness and negligence, “instructing a jury in these terms is confusing in the felony-murder context”).

<sup>160</sup> MODEL PENAL CODE § 2.02(2)(c)–(d).

substantial and unjustified.<sup>161</sup> But grading rules are rules for grading offenses: all of the actions to which they apply are already prohibited for reasons independent of the grading rule. We are permitted to impose some risks on each other because those risks are a necessary byproduct of conduct we are unwilling to prohibit: once risks become small enough, it is better simply to live with them than to try to eliminate them through criminal law. But because grading rules apply only to offenses prohibited for other reasons, it would never be better to permit such actions rather than prohibiting them, regardless of what risk they create of unintended harm. The fact that some risks are too small to warrant criminal prohibition is simply irrelevant when there already exist sufficient grounds to prohibit an action: no burglary can create a sufficiently small risk of unintended harm that it should not be prohibited, for burglary is always prohibited because it is burglary, regardless of its risk of also causing unintended harm. We must accept that we cannot eliminate all risks because the price would be eliminating all actions. But there is no tradeoff whatsoever if the price of eliminating all risk created by burglaries is eliminating all burglaries; indeed, eliminating all burglaries would be a positive good. The distinction *mens rea* draws between permissible and prohibited actions, then, is simply inapplicable for rules that distinguish among prohibited actions. Criminal law cannot bring all risk-creation within its scope, but that is no reason to ignore the risk created by actions that would in any event fall within its scope for other reasons.

Grading rules that grade offenses based on whether an offender possesses the kinds of *mens rea* criminalization rules use to define offenses, then, would not guide offenders appropriately, for the sort of mental state that makes an action prohibited is simply different than the sort of mental state that makes one offense more serious

---

<sup>161</sup> Although the Model Penal Code's definitions of recklessness and negligence clearly require that risks be both substantial and unjustified, commentators have disagreed about whether the law should in fact treat these as genuinely distinct requirements. *See, e.g.*, ALEXANDER, FERZAN, & MORSE, *supra* note 3, at 25 (rejecting the substantiality prong); Joshua Dressler, *Does One Mens Rea Fit All? Thoughts on Alexander's Unified Conception of Criminal Culpability*, 88 CALIF. L. REV. 955, 957–58 (2000) (arguing that substantiality concerns not the magnitude of the risk created but rather the degree to which it must be unjustified for creating it to be reckless); Kenneth W. Simons, *Should the Model Penal Code's Mens Rea Provisions Be Amended?*, 1 OHIO ST. J. CRIM. L. 179, 189–92 (2003) (comparing alternative views); David M. Treiman, *Recklessness and the Model Penal Code*, 9 AM. J. CRIM. L. 281, 337–38 (1981) (suggesting that the requirement of substantiality may function to exclude de minimis risks). On either approach, though, recklessness and negligence serve to distinguish risks that justify criminally prohibiting the actions that create them from risks that do not; they differ merely in whether the latter category contains only risks that are outweighed by some consideration that justifies creating them or also risks that are simply too trivial to warrant the law's attention.

than another. Applied most directly, grading rules that grade by mens rea would require factfinders to determine whether, by performing a particular offense, an offender created a substantial and unjustified risk of some further unintended harm. But offenders satisfy that standard simply by committing a crime.<sup>162</sup> For a risk to be justified, an offender's reasons for creating it must outweigh the reason against acting that the risk itself creates. If a risk is created in order to commit a crime, then, it must be unjustified: the fact that an action is a crime is not a reason to perform it but rather a reason not to perform it, and thus that reason cannot outweigh any risk, no matter how small. Furthermore, while punishing certain wrongful actions might be unreasonable if the sole reason for prohibiting them would be their creation of an insubstantial risk that were merely a de minimis wrong, crimes that create insubstantial risks are not punished merely because of the insubstantial risk they create. A primary reason to criminalize offenses in general, beyond the harm they cause intentionally, is their substantial and unjustifiable risk of causing further unintended harm. Every offender satisfies the requirements of mens rea, regardless of how risky his offense was.<sup>163</sup>

This approach evaluates the justification and substantiality of the risk an offense creates by evaluating the offense as a whole: it asks whether a risk was substantial or unjustified by considering whether the action as a whole was unjustified or substantially threatened harm. But because offenses are unjustified and substantially threaten harm—that is why they are offenses—they cannot be graded according to whether they satisfy this notion of mens rea. But perhaps grading rules

---

<sup>162</sup> Indeed, some have defended doctrines that grade by results, particularly felony murder, by arguing that offenders who commit the underlying offense are per se negligent with respect to the further harms they risk. See GUYORA BINDER, *FELONY MURDER* 27–32 (2012) (defending felony murder as a rule specifying when defendants are per se negligent with respect to the risk of death); Simons, *supra* note 23, at 1121–25 (defending felony murder and similar doctrines as forms of de facto negligence liability). Of course, if performing an offense itself satisfies a mens rea requirement, then—as these theorists recognize—grading by mens rea is not an alternative to grading by results.

<sup>163</sup> Whether the offender was reckless or negligent would depend on whether he consciously disregarded the risk or merely failed to perceive it, which would depend in turn on how the notion of conscious awareness is understood—a substantial topic in its own right, and the primary focus of chapter 2 of this dissertation. Because I will not defend an account of conscious awareness here, I will leave ambiguous whether an offender's disregard of risk should qualify as recklessness or negligence; my point is simply that some form of mens rea should apply, given that the risk disregarded, recklessly or negligently, will always qualify as substantial and unjustified. Of course, perhaps the seriousness of an offense should depend on the degree of risk disregarded only if the offender did have some relatively explicit mental state directed towards it; below, I will consider and reject a grading rule that grades offenses on that basis.



could identify offenses as more or less serious if they could evaluate the substantiality and justifiability of the risk of unintended harm considered on its own, imposing more severe sentences only on individuals who are reckless or negligent with respect to that particular risk. But it is not entirely clear how those mens rea concepts could be applied to single risks rather than entire actions. The justifiability of a risk depends on an agent's reasons for acting; but if an offender created a risk in order to commit a crime, how could the justifiability of the risk be evaluated apart from the underlying offense? Similarly, if a particular kind of offense is intrinsically dangerous—if, say, arsons often cause death—how might the magnitude of the risk be measured independent of the offense that creates it? If a risk is created by performing an offense, how could the risk be evaluated independent of the offense that creates it?

More importantly, even if risks could be evaluated independent of the offenses that create them, the distinction drawn among risks by the definitions of recklessness and negligence is not central to the function of grading rules. Grading rules identify more or less serious offenses: by sentencing riskier offenses more severely, they guide offenders to perform less risky offenses. Substantiality and unjustifiability identify the threshold separating those risks that are severe enough to justify criminal prohibition from those that are not. But while this threshold is obviously important for determining which actions are prohibited, it is not relevant when identifying which offenses are more serious than others: offenders should reduce the risk their offenses create regardless of whether the amount of risk eliminated would on its own suffice to justify prohibiting that action. Grading by recklessness or negligence would guide offenders to reduce the risk their action creates only if they can reduce it from substantial to insubstantial or from unjustifiable to justifiable (however those concepts might be applied to individual risks rather than to actions as a whole). But the guidance grading rules provide should not be limited only to this subset of risk reductions: though of course greater reductions in risk are better than smaller ones, grading rules should always guide offenders to perform less risky actions, regardless of whether a particular risk reduction would cross the threshold of justifiability or substantiality. Thus, grading by mens rea would still guide conduct ineffectively even

were it possible to use recklessness and negligence to distinguish among different offenses.

*b. Grading by Offenders' Risk Estimates*

Even if grading rules cannot guide conduct effectively by employing the mens rea concepts designed for criminalization rules, though, perhaps they could instead employ a different kind of mental state designed in particular to discharge the distinctive aims that grading rules pursue. Recklessness and negligence are ill-suited to be employed in grading rules because whether an offender satisfies their definitions depends on whether the risk his offense created would suffice on its own to criminalize conduct. By contrast, in my view grading rules should guide offenders simply to choose the offenses that they themselves estimate to be less risky. Grading rules, then, might simply grade offenses according to an offender's own awareness of risk. Such grading rules would grade offenses according to a mental state, but not by one employed by criminalization rules to define offenses: offenders who risk causing some unintended harm would receive an enhanced sentence proportional to the offender's own estimate of the magnitude of that risk.<sup>164</sup> Since offenders would expect more severe punishment for actions that they estimate to create more risk, such rules would guide offenders to perform offenses that they estimate to be less risky.

In order to employ such grading rules, however, courts would be required to actually determine how each offender estimated the risk that an offense would unintentionally create harm. And courts' efforts at determining offenders' risk estimates are likely to be even less accurate than courts' own risk estimates. To be sure, in one respect determining offenders' own risk estimates would be easier than estimating risk oneself. The latter involves two elements: courts must first identify the information about an offense to be used in estimating risk, then they must actually produce the estimate using that information. Determining offenders' own estimates, by contrast, requires only one step: courts must merely identify a particular belief that the offender held. Nonetheless, grading by offenders' risk estimates would be

---

<sup>164</sup> Alexander, Ferzan, and Morse argue that grading rules should adopt this approach. ALEXANDER, FERZAN, & MORSE, *supra* note 3, at 282.

considerably more difficult because the beliefs involved would be considerably more difficult to identify.

To grade by risk, courts must determine an offender's beliefs about the circumstances and nature of his offense. Through their behavior, offenders can provide courts with evidence that indicates which beliefs of this nature they held. Offenders typically form beliefs about the circumstances of an offense either by perception or through communication with others. And those processes each involve characteristic behaviors that can constitute evidence for the beliefs they produce: an offender's stopping to look at something is evidence that he perceived what he was looking at, just as his having a conversation with someone is evidence that he knew whichever facts his interlocutor told him. Similarly, offenders' beliefs about the nature of their actions typically depend on their intentions in acting: offenders believe that their actions will have the features they are intended to have. And offenders' intentions can ordinarily be inferred from their actions. Of course, evidence of this sort is not perfectly accurate, and it may not always be available to the courts that adjudicate offenses. Courts cannot plausibly identify all of an offender's beliefs concerning the nature and circumstances of his action. But at least it is clear what kinds of evidence courts could use in theory to determine what an offender believed.

An offender's judgments about risk, however, ordinarily are not connected closely to distinctive forms of behavior that could provide courts with evidence of those judgments. While perceptual beliefs are produced directly by an individual's interactions with the objects in the world that they are about, risk is not an object that one perceives. Instead, judgments of risk are conclusions drawn on the basis of reasoning about one's perceptual beliefs. And while interactions with objects in the world can be observed by others, and thus can constitute evidence for one's perceptual beliefs, the process of reasoning is not ordinarily observable. Consequently, courts will ordinarily lack any evidence that could indicate that an offender concluded through reasoning that an action created a particular risk. (There is an exception—reasoning is observable if it is spoken or written down—but it is unlikely that many offenders will be considerate enough to provide courts with this sort of evidence.)

Of course, to some extent an individual's factual beliefs depend on internal mental processes in addition to interactions with the world: to perceive an object, for example, one certainly must look at it, but in addition one typically must cognitively process those sensory inputs to represent them mentally and to recognize their contents. These processes, obviously, cannot be observed directly; instead, absent evidence that an offender's perceptual capacities were malfunctioning a court simply assumes that an offender saw what someone with normally functioning perceptual capacities would see. But courts cannot employ this solution to identify offenders' risk estimates. To ascribe to offenders the perceptual beliefs that individuals would normally form in particular circumstances, courts must be able to identify what beliefs those are. We can assume that someone watching a dog run across a field forms the belief that a dog is running across the field, because that's the belief that we know people would normally form. If courts could identify the correct risk estimate given the offender's information, they could perhaps ascribe that estimate to the offender. But, of course, the primary reason to rely on offenders' own risk estimates is precisely that courts cannot estimate risk correctly on their own. And if they cannot identify on their own what risk estimate would be correct or even reasonable, they cannot ascribe that estimate to the offender by assuming that he estimated risk reasonably.

Similarly, courts will ordinarily be unable to use an offender's behavior to infer how he estimated risk. At best, if the offender assigned negative utility to a particular result, courts that knew his utility function and the remainder of his credence function could establish an upper bound on his estimate of the risk of causing that result: assuming that his action had positive utility overall, his estimate of the risk of causing that result could not have been high enough so that the negative utility of that result, discounted by its likelihood, exceeded the benefits he expected from the action. But identifying all those other values on the basis of a single action is no easier than identifying the risk estimate to be sought, so however sound this approach in theory it will almost never work in practice. And in any event it establishes only an upper bound, which might be considerably higher than the offender's actual judgment of risk. (There is again an exception: whether an offender

will make wagers on the truth of a particular proposition can reveal his credence in that proposition if the utility of the wager's payoffs can be measured.<sup>165</sup> Thus, if an offender is willing to take either side of a bet at certain odds on that result's occurrence, those odds give his risk estimate. But offenders are unlikely to make a book on the results of their offense before performing it.) Ultimately, an offender's behavior is likely to be essentially the same even if his estimates of the risk that his offense will cause death differ substantially: an offender who calculates his offense creates a two percent risk of death will commit burglary or arson in the same way as one who calculates a ten percent risk, or a twenty percent risk. Because their behavior will be essentially identical, courts will have no basis upon which to distinguish them and will therefore be unable to grade their offenses differently.

Furthermore, even if evidence did exist concerning an offender's risk estimate, it would be more difficult for courts to use that evidence to determine how an offender estimated risk than to determine his factual beliefs because risk estimates are a kind of belief with a graded content. In determining whether offenders believe some matter of fact, that is, courts need only determine whether offenders possess the belief or not. Because this question can have only two answers, evidence need not be particularly precise in order to resolve it: perhaps ambiguous evidence that could support either conclusion would be difficult for courts to interpret properly, but courts should be able to reach a conclusion so long as the evidence clearly favors one over the other. By contrast, in determining an offender's risk estimate courts must identify an offender's mental state with greater precision: an offender's estimates of the risk that a particular offense creates could take any probability value, not only two. Thus, any evidence courts might use to determine an offender's risk estimates must be considerably more precise in order to differentiate among the full range of possible estimates he might have made: that evidence cannot merely justify one possible answer over the other, as it may concerning the factual beliefs that ordinarily constitute mens rea, but rather must justify ascribing to the offender a risk estimate that takes one particular value out of an entire range of possibilities. Evidence as to

---

<sup>165</sup> Frank Plumpton Ramsey, *Truth and Probability*, in *THE FOUNDATIONS OF MATHEMATICS AND OTHER LOGICAL ESSAYS* 156, 172 (R.B. Braithwaite ed., 1931).

how an offender estimated risk is generally unlikely to be available, but it is especially unlikely to be sufficiently precise to finely distinguish among the different possible values that an offender's risk estimates might take.

In addition, grading by this type of mental state presupposes that in general offenders actually will estimate the risk their offenses create. No doubt they sometimes will make such estimates concerning some results their actions may cause, particularly results that they intend to produce. But it is quite implausible that offenders will always or even often explicitly estimate the risk that an offense will produce every possible sort of harm with which the law might be concerned, particularly since many of those harms will be unintended. In reasoning, most people typically focus on achieving their own goals. Many criminals, then, will similarly reason primarily about how to achieve their criminal aims, rather than bothering to explicitly consider the likelihood that a particular offense will cause harms unrelated to those aims. The possibility that offenders will simply fail to estimate risk poses a further practical challenge for courts: just no evidence is likely to distinguish between arsonists who differ in how they estimate the risk of causing death but who each perform an identical crime, it is unlikely to distinguish an arsonist who performs the same crime without thinking about the risk of death at all. But such offenders also pose a theoretical problem for grading by this sort of mental state: if offenders should be sentenced according to their own risk estimates, what sentence should be imposed on an offender who lacks any risk estimate? Since the offender lacked any risk estimate and since courts cannot estimate risk accurately, it is not clear how any risk estimate could be used in sentencing such offenders. But imposing no additional sentence would provide counterproductive guidance: since such grading rules would impose the least severe sentences on those who did not estimate the risk their offenses would create, they would guide offenders not to perform the offense they estimate to be least likely to cause harm but rather to avoid estimating the risk created by any of their options at all and instead to choose among them on some other basis.<sup>166</sup>

---

<sup>166</sup> Alexander, Ferzan, and Morse, the leading defenders of grading rules that grade by an offender's own risk estimates, defend this approach: "After a jury determines which legally protected interests the actor believed himself to be risking, the jury will need to discount these interests by the actor's belief as to the magnitudes of the various risks he was imposing." ALEXANDER, FERZAN, & MORSE, *supra* note 3, at 282. It is worth noting that even their approach does not fully solve the problem of offenders who failed to estimate risk. If an offender believes

Finally, even if offenders were to explicitly consider the risks that different offenses create, it is not plausible that they will produce estimates of those risks in a form that courts could use in sentencing. For to the extent that offenders do consider those risks, they are likely to make only comparative judgments about them—that is, only to determine that one offense would create a greater or lower risk than another. But these judgments alone would not suffice for courts to impose sentences—how should a court sentence an offender who performed an offense in a manner that he judged was riskier than some alternatives but less risky than others? Rather, just as when courts themselves estimate risk, offenders could make usable risk estimates only by using some standardized scale that measures all risks relative to certainty and impossibility and that courts could therefore use to translate risks into a common sentencing scheme.<sup>167</sup> But for precisely the same reasons that courts cannot accurately make these estimates, offenders are unlikely to make them accurately in this form, either: offenders cannot enumerate all possibilities compatible with their information and thus cannot compute the fraction of such possibilities in which harm occurs.<sup>168</sup> And, psychologically speaking, it is quite implausible that agents actually ever make such judgments: when choosing between multiple possible actions it is important to compare them to one another, but there is no reason why one would need to compare the risk created by any of those actions to the risk created by any other action not under consideration.<sup>169</sup> Even when offenders do explicitly compare risks, then, they

---

that his offense creates a particular risk but gives no thought to the precise degree of risk it created, their approach appears to require that risk to influence his sentence, since he did believe himself to be creating it, but it determines the extent of the influence based on something—his belief as to the magnitude of that risk—that simply does not exist. This problem strikes me as quite pressing: I suspect few offenders would deny that they recognized the existence of certain obvious risks but that many could quite honestly say they did not bother to explicitly calculate their magnitude.

<sup>167</sup> See *supra* note 145 and accompanying text.

<sup>168</sup> See *id.*

<sup>169</sup> Alexander, Ferzan, and Morse appear to think that this problem could be solved by employing a scale with relatively few risk categories: they suggest that since individuals probably do not “think in numeric probabilities,” it would be more psychologically realistic to grade all offenses on a scale with six categories ranging from virtual certainty to virtually no risk. ALEXANDER, FERZAN, & MORSE, *supra* note 3, at 282. But since they provide a translation between their six-part scale and percentages—i.e. “low risk” covers percentages ranging between twenty and forty—it is deeply unclear how their proposal achieves psychological realism. If their translation scheme is accurate, then individuals estimate risk to be low if and only if they estimate it to be between twenty and forty percent. This degree of precision would seem miraculous unless individuals were thinking about risk in terms of percentages. More generally, though, to distinguish a six-part scale from a numerical one misunderstands what numeric probabilities are: a scale ranging from zero to five employs fewer numbers than a scale ranging from zero to one hundred, but it is hard to see how reasoning ceases to become numeric simply because fewer numbers are used, or if they are expressed in forms like “3/5” and “5/5” rather than “60%” and “100%.” Rather, as Keynes argued, what makes probabilistic reasoning numeric is that probabilities are measured on some common scale

will not produce the explicit risk estimates that courts would require in order to grade offenses by the offender's mental states.

As a purely practical matter, then, offenders' own risk estimates are an impractical basis with which to grade their offenses. In addition, though, grading by those estimates would not guide offenders' conduct effectively even in theory. I suggested that grading by mental state might be an effective way of guiding offenders based on an analogy to the kinds of mens rea that criminalization rules employ in defining intentional crimes. But there is a key dissimilarity between intended and unintended results that prevents grading by mental state from effectively guiding offenders to reduce the risk that their offenses will unintentionally cause harm. Given that individuals are by and large effective at acting to achieve their aims, actions intended to produce a harmful result are an especially dangerous kind of behavior—actions intended to cause a harmful result normally do cause it. Because these actions are dangerous in virtue of the intention with which they are performed, the law properly forbids individuals from acting with such intentions, thereby guiding them instead to act with other kinds of intentions. Since the intention involved in such actions itself explains why they are so dangerous, guiding offenders to act with other intentions in itself eliminates the danger those actions pose. Of course, because actions intended to produce a particular harmful result are not the only kind of action likely enough to produce harm for criminal prohibitions to be justified, there exist other offenses defined through types of mens rea other than the intention to produce a particular harm. But because actions are likely to cause harm if they are performed

---

applicable to all probabilities. *See generally* KEYNES, *supra* note 129, at 182 (showing how numerically measurable probabilities exist so long as they can be placed on a scale ranging from certainty to impossibility, which may have any number of values in between those two endpoints). And it is psychologically unrealistic to suppose that probabilistic reasoning ordinarily satisfies this standard: individuals do typically assess risks in terms such as "low" and "high," but those judgments that a risk is "low" or "high" do not have a constant meaning in all contexts. The very same action may be judged low risk or high risk depending on what it is being compared to—burning down an empty warehouse is low risk if the alternative is burning down an occupied dwelling, but high risk if the alternative is burning a pile of leaves in one's yard. Similarly, two different actions may each be judged low risk, given the alternatives, even if the agent forms no judgment as to their risk with respect to one another—one may think that burning down an empty warehouse is low risk, since the alternative is burning down an occupied dwelling, and that selling pure heroin is low risk, since the alternative is selling fentanyl-laced heroin, without having formed any judgment about whether burning down an empty warehouse or selling pure heroin is higher risk. But if the meanings of low risk or high risk are not constant, how could the law translate them into a constant sentencing scheme?



with the intention to cause it, the law can reduce the occurrence of harm by requiring individuals not to act on that intention.

But although intentional wrongdoing typically produces the harm intended precisely because the offender intended to produce it, offenses do not unintentionally cause harm because an offender estimated that they had a high risk of causing it. That is, although offenders are well placed to accurately estimate the risk their offenses create, a given offense would be just as likely to cause harm if an offender estimated its risk differently or failed to estimate its risk entirely. Because intentions guide action, the law can change offenders' actions by changing the intentions with which they act: the intention to produce a particular harm guides an offender in actually producing it, and thus requiring individuals not to act on that intention makes them less likely to produce that harm. But a risk estimate on its own need not guide behavior at all—that estimate is simply a belief, not an intention. Thus, while grading rules that punish offenders based on their risk estimates will guide them to avoid an action that they believe will create a high risk of harm, they need not do so by choosing a different action—they may instead adopt a different belief, either by avoiding explicit reasoning about risk entirely or by estimating risk using methods that tend to underestimate it. And since the offender's estimate does not causally affect the risk an action creates in the way that an intention causally affects the action an individual performs, changing how an offender estimated risk will not itself make his offense any less risky in the way that changing his intentions will. Because risk estimates need not causally control behavior, grading rules cannot simply presuppose that offenders will choose offenses using their estimates of the risk that different possible offenses would; instead, such rules must first actively induce offenders to make accurate estimates concerning unintended results, then guide them to use those estimates as a basis for choice. And while grading by subjective risk estimate does guide offenders to choose the offense they estimate to be least risky once they have estimated risk, it does not instruct them to make the estimates in the first place.

Ultimately, then, no form of grading by mental state effectively guides offenders to reduce the risk their actions create. The kinds of mens rea employed by criminalization rules are concerned primarily with a specific threshold between risks,

one that separates risks that do and do not justify prohibiting the action creating them. But this threshold is largely irrelevant to the function grading rules perform—namely, guiding offenders to perform actions that are less risky. And while simply grading offenses by the offender’s own risk estimate would not assign undue importance to a threshold relevant only for criminalization questions, such grading rules would be ineffective both practically and theoretically. Practically, because risk estimates are the result of an inferential process that offenders will likely perform only mentally, if at all, it is not clear what evidence could even be used to identify how an offender estimated risk, if he estimated it at all. Theoretically, because offenders do not need to produce risk estimates or employ them in deliberation and choice, grading rules must influence offenders first actually to produce those estimates, and only then to prefer offenses whose risk they estimate to be lower. But grading by mental state only provides the latter sort of guidance. If offenders do estimate the risk of two possible offenses, then since they should expect to be punished more severely if they perform the one they estimate to be riskier, grading by mental state guides them to prefer the other. But if they do not estimate the risk of either action, grading by risk does not guide their action at all, even if one is likelier to cause harm than the other; indeed, if risks that they do not estimate simply do not influence their sentence, grading by mental state would encourage offenders to ignore the risk of unintended harm entirely, the exact opposite of the guidance it should provide. Thus, it does not effectively guide offenders to act based on their own risk estimates.

## *2. Grading by Results*

In order to guide offenders to choose among offenses based on their own estimates of risk, then, criminal law must employ some alternative method of grading. In particular, grading by mental state fails to effectively guide offenders to reduce the risks their offenses create because the harms risked are unintended. Offenders must actually estimate the risk their offenses would create in order to choose how to offend on the basis of those estimates. But while grading by mental state might be able to effectively guide offenders in choosing among different offenses if those estimates do exist, an offender who intends neither to cause nor to avoid those harms will ordinarily neither estimate how likely they are to occur nor choose on the basis of

those estimates. Furthermore, precisely because those harms are not intended, offenders' attitudes towards them will not ordinarily be among the mental states that influence their conduct, and thus courts will struggle to infer how an offender estimated that risk from the manner in which the offense occurred. Thus, grading rules cannot achieve their function simply by instructing offenders not to act with a particular intention, as criminalization rules can; rather, grading rules must guide offenders to act with a particular intention—namely, the intention of reducing the risk that their offense will cause harm. For if offenders intend to reduce the risk that their offenses will cause harm, that intention will guide them automatically to estimate that risk and use those estimates in deliberation. How, then, can grading rules induce offenders to intentionally reduce the risk their offenses create?

In general, grading rules guide conduct because individuals will choose actions that the law punishes less severely (whether because they accept the judgments of seriousness those punishments express or because they respond to the incentives those punishments create). Thus, grading rules can induce changes in offenders' intentions through how they punish different offenses: if grading rules punish one kind of offense more severely than another, they will thereby induce offenders to perform the latter kind of offense rather than the former. One way of providing that guidance, of course, would be to grade offenses by their risk. If riskier offenses are punished more severely, then an offender who wants to reduce the scope of his potential criminal liability must intentionally choose an offense that creates less risk. Since such rules grade offenses according to the law's risk estimates, not according to offenders' own estimates, they do induce offenders to actually estimate risk rather than guiding conduct only if offenders have already estimated risk. But, as I have argued, such rules guide offenders ineffectively because they guide offenders to employ the wrong risk estimates. Grading rules that grade by risk grade, in particular, by the law's estimates of risk; because all grading rules guide offenders to avoid actions that they believe the law will punish more severely, grading by risk guides offenders to perform offenses that they believe the law will estimate to be less risky. But because the law will often estimate risk inaccurately, such rules will guide offenders incorrectly: they will choose offenses whose risk the law estimates to be

low even if those offenses are not actually lowest in risk. Because individuals are ultimately guided by their expectations of how courts will apply grading rules, courts cannot guide offenders to employ accurate risk estimates unless courts can estimate risk accurately, which they cannot reliably do.

Instructing offenders to perform the action that has the lowest risk of causing harm might seem like the most natural way to guide offenders so that they avoid unintended harms: if risk is what they ought to reduce, the law can simply tell them to reduce risk. But in some ways this guidance is in fact atypical. For we do not ordinarily guide others' conduct by telling them which epistemic states they must have towards the actions they choose: we would tell an unpunctual friend not to be late rather than telling him not to perform actions that he believes will likely cause him to be late, just as we might advise a student to choose the classes taught by the best professors, not the classes that she believes will likely be taught by the best professors. Ordinarily, that is, we simply tell people which results to pursue, then leave up to them to determine which actions will likely achieve those results. To be sure, I have argued that criminalization rules cannot pursue that strategy: because criminalization rules provide decisive criteria for deliberation, they must directly identify which actions individuals must exclude from deliberation. But grading rules provide non-decisive criteria, which influence deliberation without dictating its results, so they may allow offenders themselves ultimately to translate that criterion into choice. Thus, grading rules can simply instruct offenders to avoid causing harm, leaving offenders themselves to determine which offenses are least likely to cause it.

Such rules, of course, would grade by results: punishing an offense more severely if it causes a particular harm communicates to offenders that they should intentionally avoid causing those harms. Instead of inducing the intention to perform the action that is least risky, such rules simply induce offenders to intend not to cause harm—or, at least, to intend to perform the underlying offense without causing any additional harm. Precisely because such rules identify only the harm to be avoided without further identifying which actions would effectively avoid it, offenders themselves must evaluate which of their options would best implement their intention to avoid harm. And to do so, of course, they must estimate the risk of harm that each

option would create. By sentencing offenders based on the results of their offenses, the law guides them to choose among offenses based on their likelihood of causing harm.

In general, offenders face substantial uncertainty concerning the unintended harms that their offenses cause: though offenders possess some information relevant to determining whether an offense will cause a particular result, much relevant information will remain out of their grasp. Nonetheless, despite this uncertainty unintended harms are a subject of considerable importance to the criminal law: even though it is uncertain which particular offense will cause harm, offenders ought in general to perform offenses in a manner that reduces the harm they cause in general. Grading by risk attempts to eliminate this uncertainty from offenders' decision-making. Each of its different forms proposes some method of measuring uncertainty based on offenders' own information, whether a phalanx of rules defining risk factors, a procedure for courts to estimate likelihoods directly, or a requirement that courts defer to offenders' own beliefs about risk. Then, grading by risk instructs offenders to reduce risk, thus measured. In a sense, then, grading by risk replaces decision-making under imperfect information with a new form of decision-making that offenders can implement using only their actual information. Since grading by risk requires offenders to reduce risk as judged by one of various metrics applied to the information they possess, they can decide based simply on the direct application of that metric to their actual information. Though offenders do not possess all the information that determines whether a particular result will occur, grading by risk measures risk, and thereby guides choice, using their actual information. Thus, though offenders have only incomplete information with respect to which offenses cause harm, they have all the information necessary for the decision-making procedure that grading by risk requires them to perform.

But while replacing a decision made under imperfect information with a different decision made under full information would be valuable, could it be done correctly, none of the metrics grading by risk proposes for measuring risk are very good. Thus, the law cannot give individuals a rule to regulate their conduct that depends only on information they possess; instead, the decisions grading rules

regulate must be made under imperfect information. Unlike grading by risk, which attempts to devise a method of decision-making under perfect information to replace the decision-making under imperfect information that individuals actually confront, grading by results accommodates the uncertainty that offenders face by introducing uncertainty into the guidance it provides. Grading rules guide offenders to choose offenses based on their judgments as to which offenses are least likely to cause harm. They exert that influence through the severity of the sentences they impose, since offenders will prefer offenses that are punished less severely. Thus, grading rules can guide offenders to perform the offenses they judge to be least risky if offenders expect those offenses to be punished least severely. But rather than employing a grading rule that directly identifies which actions those are, which would require the law somehow to identify which actions are least risky, the law can produce the same guidance if offenders have the same uncertainty about the results of actions as about the punishments to be expected for performing them. That is, offenders might expect the least risky action to receive the least severe sentence not only if some rule explicitly requires that action to be sentenced least severely but also if an action's risk of receiving a more severe sentence is the same as its risk of causing a particular harm. Such an offender would face uncertainty about which offenses would be punished more severely, of course, but so long as he faces the same uncertainty about punishment as about harmful results, his uncertainty about punishment enables the law to guide conduct effectively rather than undermining it. And the surest way to make the risk of punishment automatically match the risk of causing harm is to impose a more severe punishment only if an offense causes harm—in short, to grade by results.<sup>170</sup>

Grading by results, then, successfully guides offenders to choose offenses based on their own estimates of the risk those offenses create. As I have noted, estimating risk involves two components—first identifying the information to be used as the basis for the estimate, then actually producing it. Because offenders face considerably fewer obstacles than government entities in performing each component,

---

<sup>170</sup> This observation lies at the core of David Lewis's defense of imposing higher sentences for successful than for unsuccessful attempts. See Lewis, *supra* note 17, at 63–67.

their estimates are likely to be more accurate than estimates produced by law; consequently, the risk estimates offenders themselves produce will more effectively guide them to avoid offenses that cause harm. First, in order to estimate the risk that an offense creates, a court must identify those facts about how it was performed that will be used to determine its risk. But as I have argued, courts are unlikely to correctly identify every fact about an offense relevant to the risk that it creates. And because facts that courts do not identify do not affect the sentences they impose, grading by risk does not guide offenders to reduce the risk of offending except with respect to facts that they expect courts to be able to successfully identify. Grading by result does not limit its guidance in this way. Since grading by result increases the punishment to be imposed on an offender just in case his offense actually causes harm, such rules simply guide offenders not to cause that harm. That sort of guidance draws no distinctions among the different ways in which offenders might reduce the likelihood of causing that harm by changing the manner in which they perform the offense. Instead, because any measure that might affect the likelihood that the offense causes harm might affect how it is graded, all such measures are automatically relevant to an offender's deliberation. Grading by result guides offenders to reduce the risk their offenses create in any manner that actually will affect its risk. So long as offenders can identify some possible causal effect, it does not matter whether courts or anyone else are able to identify how the offense produced that effect.

Second, because grading by risk imposes sentences based on how some government entity estimates the risk a particular offense created, it guides offenders to choose the offense that they believe the government will estimate to be least risky. But in order to estimate risk any government entity must employ some methodology for measuring risk: the facts of each offense must somehow be quantified and translated into an actual sentence, measured in the months of a prison sentence or the dollars of a fine. And since the law must sentence every crime, grading by risk requires some set of grading rules that courts can use to quantify the risk any offense creates: those rules must apply universally, to all offenses, and therefore must implicitly or explicitly compare all criminal offenses with respect to their risk. One approach, grading by risk factor, makes those comparisons explicit. It identifies in

advance the extent to which every aspect of how an offense might be performed increases or decreases the risk that it creates; those rules collectively quantify how any possible offense would compare in risk to any other. The other approach, grading directly by risk, makes those comparisons implicitly. It instructs each court imposing a sentence only to directly estimate the risk the offense created. But because grading directly by risk must translate all those risk estimates into sentences, it must employ some common scale for doing so—one ranging from the certainty of harm to the impossibility of harm. Thus, since grading directly by risk uses a common scale to measure the risk created by each offense, it implicitly compares the risks created by all possible offenses to one another: by comparing each offense's likelihood of causing harm with certainty, grading directly by risk indirectly produces comparisons of the risk created by all offenses that are sentenced.

But estimating the risk individual offenses create through a scheme for comparing the risks created by all possible offenses possible offense is a considerably more ambitious project than is necessary to guide individuals' conduct. For an offender deliberating about how to perform an offense, it does not particularly matter how the likelihood that one particular option will cause harm compares to certainty, or to impossibility, or to most other possible actions. Since his choice is limited to the options between which he is deliberating, the only comparison relevant to his decision is a comparison between the risks that those options create. Certainly, those options could be compared in risk with all possible actions. In particular, grading by risk guides each offender's choice among his options by ranking each according to some universal system that compares it to all possible offenses. But since the offender cannot choose any possible offense, his actual deliberation does not depend on that broader comparison. Grading by risk attempts a more ambitious task than is necessary to guide conduct appropriately. And that more ambitious task—comparing the risk of all offenses—is considerably more difficult than the task of comparing the small subset of actions between which an offender must choose: it is considerably easier to define a partial ordering stating that some probabilities are higher or lower than others than it is to define a complete ordering enabling any two probabilities whatsoever to



be compared.<sup>171</sup> As I have argued, it is simply not plausible that sentencing law can accurately compare the risk created by every possible offense. The complex systems of universal risk estimation grading by risk employs will sometimes judge risk inaccurately, and because grading by risk guides offenders in comparing their options only by placing those offenses on a universal scale that completely ranks all risks, inaccuracies in the universal scale employed will produce incorrect guidance for how offenders should choose among their options.

Grading by result, however, completely avoids the problem of producing some universal system for measuring risk. Under grading by result, the law expresses judgments only about which results are harmful and must be avoided. Offenders themselves, then, must determine which of their options is likeliest to avoid the harms that grading rules require them to avoid. Because offenders themselves face a choice only among some limited set of options, they may simply compare those options directly: they have no reason to try to measure those options using a scale that enables comparisons to all possible offenses. And since it is easier to make comparative probability judgments among a small number of actions than to define some universal scale that ranks the risks created by any action, offenders' own comparisons among their options are likely to guide their conduct more effectively than the law's methodology for explicitly estimating risk. Sometimes, that is, an offender will be able to recognize that one possible offense is riskier than another, even though the law's method for measuring the risks created by all offenses will introduce inaccuracies in how it measures the risks those offenses create. Since grading by risk guides offenders to act according to those inaccurate estimates, it will reduce harm less effectively than grading by result, which instructs offenders to act according to their own judgments as to which offenses are riskier than the alternatives.

In short, grading by risk guides conduct ineffectively because it relies on the law to produce and employ some model for estimating risk. Because courts sentence offenders based on that model of how offenses create risk, grading by risk guides offenders to reduce risk as measured by that method. But because of errors both in the

---

<sup>171</sup> See generally KEYNES, *supra* note 129, at 37–42 (explaining how probabilities in general should be understood as partial orderings, and arguing that only in rare and special cases can those partial orderings be understood as representing numerical probabilities).

model's inputs—the imperfect information about each offense that courts use to estimate its risk—and in its structure—the inaccuracies of the universal scale measuring all actions' risk that it unnecessarily produces—grading by risk ultimately guides conduct ineffectively. It fails to guide offenders to reduce risk in ways that are not represented by the model's inputs or are not measured properly by its methods of estimating risk. Grading by result, instead, guides conduct without employing any model for estimating risk: it relies on the actual causal relations in the world rather than on some theoretical model of how those relations function. Because sentences depend only on the results of conduct, offenders deliberating over how to offend must consider only whether a particular offense actually will cause a result, not how a court will judge its risk of causing that result. And while the guidance provided by a model is limited to the factors incorporated into the model—offenders need only consider aspects of offenses that the model uses to estimate risk and that it can measure properly—there is no limit on what information is relevant if offenders are judged simply by the results of their actions. For if results themselves matter, offenders must do whatever they can in trying to produce the required result: they must consider every feature of an offense that might affect whether it causes that result, not merely those used as inputs to some model, and they must pick the offense they believe actually to be least likely to cause that result, regardless of how each offense is measured on some universal scale for estimating risk. Since results matter, offenders will be guided by their actual uncertainty as to the results of their actions, not by an explicit method of reconstructing and quantifying that uncertainty.

Ultimately, then, whether grading by result or grading by risk will guide conduct more effectively depends on a comparison of two very different kinds of individual capacities. Grading rules explicitly articulate in language the principles offenders must employ in deliberating over how to offend. Under grading by risk, those rules state how to judge the risk individual offenses create. By contrast, under grading by result, those rules state only which results offenders are to avoid. On this approach, no explicit rules state which offenses are least likely to produce those results; offenders must simply avoid those results in action. While grading by risk depends on the linguistic capacity to articulate explicit rules for how to avoid harm,

grading by result depends on the simple capacity to avoid harm in acting. Thus, the relative success of the two depends on whether our capacities to act effectively outrun our capacities to explicitly articulate verbal rules for acting effectively. If the former, the law will guide action more effectively through rules that simply instruct individuals to avoid certain harms, then rely on their capacities and skills as ordinary agents to actually avoid it; if the latter, the law will guide action more effectively through the explicit grading rules it institutes to identify which actions are riskier than others.

But at least concerning activities that individuals regularly perform, it is in general easier to successfully execute a complex task than it is to give explicit rules for successfully performing that task. It is easier to ride a bike than to give explicit rules stating the precise movements of one's muscles necessary to ride a bike; it is easier to speak a language grammatically than it is to give explicit rules of grammar; it is easier to tell a funny joke than it is to give explicit rules governing humor.<sup>172</sup> In my criticisms of grading by risk, I have suggested that the same is true of the sorts of risks with which grading rules are concerned: it is easier to interact with physical objects in a dangerous manner without causing harm, for example, than it is to give explicit rules stating how to conduct those interactions without causing harm.<sup>173</sup> If

---

<sup>172</sup> GILBERT RYLE, *THE CONCEPT OF MIND* 30 (discussing “performances in which intelligence is displayed, but the rules or criteria of which are unformulated”); Charles Wallis, *Consciousness, Context, and Know-How*, 160 *SYNTHESE* 123, 130–38 (2008) (summarizing empirical research in psychology and cognitive science to conclude that the capacity to successfully perform complex tasks does not generally require the ability to verbally represent one's procedures for performing those tasks). Ryle himself took the phenomenon as evidence for the independence of “knowing how” from “knowing that”—namely, that the ability to perform a complex task or to exercise some skill cannot be understood as a relation between the individual executing the performance or the skill and some proposition that she could be said to know. That claim has been controversial; other philosophers have argued that knowing how is really a special kind of knowing that, a position commonly known as intellectualism. But regardless of whether Ryle or his critics are correct on that point, the kinds of arguments used to defend intellectualism concede the conclusion I wish to draw from these examples. For contemporary intellectualists typically acknowledge that the sort of knowing that upon which knowing how depends cannot readily be explicitly expressed in words. *E.g.*, Jerry A. Fodor, *The Appeal to Tacit Knowledge in Psychological Explanation*, 65 *THE JOURNAL OF PHILOSOPHY* 627, 634 (1968) (“Certain of the anti-intellectualist arguments fail to go through because they confuse knowing that with being able to explain how.”); JASON STANLEY, *KNOW HOW* 162 (2011) (“[I]t is thoroughly unclear why it should be a condition on propositional knowledge that knowing that p requires that one possess an accurate informative verbal description, in the required sense, of the content that p.”). Whatever may be true of knowledge, for the law to effectively guide individuals to act in a certain way, it must possess an “accurate informative verbal description.” Thus, even if individuals' ability to perform complex skills by following rules they cannot explicitly articulate is not good evidence that their capacities fail to reduce to propositional knowledge, it is certainly good reason to think that explicit legal rules could not effectively guide them in exercising those capacities.

<sup>173</sup> Wallis, *supra* note 172, at 138 (“Knowledge of how to perform physical skills is often acquired without involving areas of the brain associated with conscious executive functioning. Indeed, knowledge of how to

that is correct, then grading by result will ultimately be a more effective way of guiding offenders to reduce the harm their offenses create than any set of grading rules instructing offenders to reduce risk as measured by some explicit procedure for estimating risk: it is better to rely on offenders' capacities for skillful action than to give them a flawed set of explicit rules they must use to govern their choices. To be sure, I have agreed that some explicit rules may helpfully instruct individuals how to reduce the harm their offenses create; thus, grading rules may effectively give offenders some rules to rely upon in acting when those rules are accurate. But because accurately devising such rules will be difficult, possible offenses will differ in many ways that do influence the risk they create but that are not reflected in any method of explicitly estimating risk. Consequently, the law will most effectively reduce the harm offenses cause by relying in large part on individuals themselves to determine what manner of offending is least likely to cause harm—that is, by grading by results.

### **III. Conclusion**

This chapter has defended rules of sentencing law that grade the seriousness of offenses based on the harms they cause. My argument has depended largely on an understanding of the function of those rules—namely, to guide offenders to offend in a manner less likely to cause harm. Criminal law should grade offenses by their unintended results, in my view, because such grading rules most effectively discharge that function. The function of grading rules undergirds the distinction between grading rules and criminalization rules, which ordinarily do not define offenses in terms of results alone: this difference in their functions explains why grading rules may employ results alone even though criminalization rules must define the minimum conditions of criminal liability in terms of a mental state of the offender. And my comparison between grading rules that grade in various ways by risk and those that grade by results focused on which approach could more effectively identify which offenses offenders should choose to perform: because producing accurate risk estimates is so difficult, courts can better guide conduct by simply identifying the

---

perform skills often literally moves from areas of the brain associated with conscious executive functioning to areas of the brain strongly disassociated with conscious executive functioning.”).

results individuals must aim not to produce without further specifying how individuals must aim to reduce the risk of producing them. Opponents of grading by result correctly note that differences in the results of criminal conduct need not on their own indicate some difference in the offender's disregard of the interests of others. But criminal liability is not imposed merely to record the extent of that disregard; criminal liability communicates rules to govern individuals' behavior. Thus, principles governing the imposition of liability must consider how the success of that communication depends on what principles govern the imposition of liability. And, in my view, grading most successfully communicates to offenders how to act if the results of criminal activity influence the extent of a defendant's liability.

## Chapter 2: Recklessness and the Awareness of Risk

Paradigmatic instances of crime involve the intentional performance of an action that satisfies the elements of a criminal offense—intentionally injuring another, say, or intentionally taking his property. And intentionally performing a particular type of action involves acting with a mental state, such as intention, purpose, or knowledge, directed towards the features that define actions of that type. To be guilty of such an offense, then, a defendant must possess the mental state the offense requires directed towards at least some of the elements that define it.<sup>1</sup> Not all offenses concern intentional wrongdoing, though; other grades of mens rea define unintentional crimes that defendants may commit without an intention directed towards the elements of any offense. To the extent that culpability for such offenses depends on the defendant's mental states,<sup>2</sup> it must extend culpability beyond the mental states involved in intentional wrongdoing.<sup>3</sup> And while the mental state required to engage in intentional wrongdoing is directly identified by its definition—wrongdoing is intentional because of the involvement of some intention or related mental state—the definition of unintentional wrongdoing does not so clearly identify any mental state required for culpability, since merely specifying that a particular kind of intention is not required obviously does not identify what alternative mental state must be involved.

---

<sup>1</sup> Only some because mens rea need not always be required for each element of an offense.

<sup>2</sup> Of course, the culpability of unintentional wrongdoing might not require a particular a mental state—as, on some analyses, criminal negligence does not. See, e.g., Michael S. Moore & Heidi M. Hurd, *Punishing the Awkward, the Stupid, the Weak, and the Selfish: The Culpability of Negligence*, 5 CRIM. L. & PHIL. 147, 148–50 (2011).

<sup>3</sup> Cf. R.A. DUFF, INTENTION, AGENCY, AND CRIMINAL LIABILITY: PHILOSOPHY OF ACTION AND THE CRIMINAL LAW 139–42 (1990) (analyzing recklessness as an extension of the paradigm of intentional wrongdoing).

Standard legal definitions of recklessness are similarly unclear. For example, the Model Penal Code’s definition, the template for many state codes,<sup>4</sup> requires a defendant to “consciously disregard a substantial and unjustifiable risk.”<sup>5</sup> But like the notion of unintentional wrongdoing itself, this definition characterizes recklessness more by what it does not involve than by what it does. To disregard a risk is simply not to give it an appropriate role—not to regard it, or at least not to regard it adequately—in deliberating about what to do. But all agents who perform risky actions do not assign that risk the appropriate role in deliberation, including those who act purely by accident or from ignorance, for had they regarded it appropriately they would not have acted. Thus, disregard of risk alone cannot always be culpable. And the Model Penal Code’s proposed basis for isolating culpable disregard—namely, whether or not it is conscious—is singularly unhelpful: the adjective merely suggests that the disregard must involve some conscious mental state while providing virtually no guidance as to which conscious mental state that must be. Further ambiguity is added by the generic reference to “risk” with no account of exactly what a risk is, leaving unclear exactly what a reckless defendant must disregard. The definition is ambiguous concerning both the kind of mental state involved in recklessness and its object. And the ambiguity of legal definitions of recklessness extends beyond the Model Penal Code: as Findlay Stark has documented, “the main Anglo-American systems of criminal law have failed to explain the concept of awareness of risk in any depth.”<sup>6</sup> Given the importance of that concept to the definition of recklessness, Douglas Husak has described this failure as “scandalous.”<sup>7</sup>

This chapter will attempt to address that failure: I will present and defend an account of the mental state involved in recklessness. Given the clear role that risk plays in the notion of recklessness, some theorists have analyzed it by extending the paradigm of intentional wrongdoing to incorporate the notion of risk. In particular, if intentional wrongdoing involves an intention to perform an action with certain features—features that do, in fact, satisfy the elements of some offense—then

---

<sup>4</sup> See David M. Treiman, *Recklessness and the Model Penal Code*, 9 AM. J. CRIM. L. 281, 375–86 (1981).

<sup>5</sup> MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962).

<sup>6</sup> FINDLAY STARK, *CULPABLE CARELESSNESS: RECKLESSNESS AND NEGLIGENCE IN CRIMINAL LAW* 90 (2016).

<sup>7</sup> Douglas Husak, *Negligence, Belief, Blame and Criminal Liability: The Special Case of Forgetting*, 5 CRIM. L. & PHIL. 199, 207 (2011).

recklessness might be understood as requiring an intention to risk performing an action that has those features. According to this approach, recklessness is intentional risk-taking, or choosing to create a risk.<sup>8</sup> Whether a risk may be taken depends on its degree—whether, in the Model Penal Code’s words, it is “substantial and unjustifiable.”<sup>9</sup> And the consciousness or awareness of risk required for recklessness then follows from the general analysis of intentional wrongdoing. The sort of awareness required for culpability for intentional wrongdoing is satisfied by a mental state—an intention or belief—directed towards the features that make that action wrongful. If recklessness instead involves awareness of risk, then, it should require a similar mental state directed towards the risk that an action will have those features.<sup>10</sup>

In this chapter, I will reject the idea that some belief about the existence of risk itself is necessary for a defendant to be reckless. To be sure, accounts of recklessness do exist that deny any such mental state is necessary for a defendant to be reckless.<sup>11</sup> In rejecting interpretations of recklessness as intentional risk-creation, however, these accounts do not dispute what mental state would be required for recklessness were it to require some culpable mental state directed towards risk: they generally agree that the awareness of risk should be understood to require some belief

---

<sup>8</sup> For approaches along these lines, see, for example, LARRY ALEXANDER, KIMBERLY KESSLER FERZAN & STEPHEN J. MORSE, *CRIME AND CULPABILITY: A THEORY OF CRIMINAL LAW* 192 (2009) (“[C]ulpability should attach to those acts of knowingly taking unjustifiable risks . . . .”); DUFF, *supra* note 3, at 142 (critically discussing “the view that recklessness should be defined as conscious risk-taking”); STARK, *supra* note 6, at 141 (“Without a belief about the relevant risk, no relevant choice to take it is made.”); Kimberly Kessler Ferzan, *Opaque Recklessness*, 91 J. CRIM. L. & CRIMINOLOGY 597, 631 (2001) (describing an individual as reckless when he “consciously chooses to engage in risky behavior”); Jenny McEwan & St. John Robilliard, *Recklessness: The House of Lords and the Criminal Law*, 1 LEGAL STUD. 267, 268 (1981) (“Both the leading textbooks [of English law] describe recklessness as deliberate risk-taking.”); Glanville Williams, *The Unresolved Problem of Recklessness*, 8 LEGAL STUD. 74, 90 (1988) (describing reckless defendants as “consciously decid[ing] to run a risk”).

<sup>9</sup> MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962).

<sup>10</sup> *E.g.*, ALEXANDER, FERZAN, & MORSE, *supra* note 8, at 282 (“After a jury determines which legally protected interests the actor believed himself to be risking, the jury will need to discount these interests by the actor’s belief as to the magnitudes of the various risks he was imposing.”); STARK, *supra* note 6, at 91 (“[T]o be aware of a particular risk associated with  $\phi$ -ing requires a belief that the relevant risk exists.”); Ferzan, *supra* note 8, at 600 (arguing that reckless agents “know the dangerousness of their acts and therefore can decide whether to commit those acts”); Husak, *supra* note 7, at 208 (“[T]he reckless defendant . . . believes that he is creating a substantial and unjustifiable risk.”); Alexander Sarch, *Review of Findlay Stark, Culpable Carelessness: Recklessness and Negligence in the Criminal Law*, 12 CRIM. L. & PHIL. 725, 730 (2018) (“[I]t seems preferable to construe awareness of a risk as the true belief that the risk exists.”). *But see* Marcia Baron, *Negligence, Mens Rea, and What We Want the Element of Mens Rea to Provide*, 14 CRIM. L. & PHIL. 69, 71 n.7 (2020) (disputing whether recklessness should be understood in terms of beliefs).

<sup>11</sup> *See, e.g.*, DUFF, *supra* note 3, at 139–79; VICTOR TADROS, *CRIMINAL RESPONSIBILITY* 237–64 (2007); Alan C. Michaels, *Acceptance: The Missing Mental State*, 71 S. CAL. L. REV. 953 (1998); Samuel H. Pillsbury, *Crimes of Indifference*, 49 RUTGERS L. REV. 105 (1996); Kenneth W. Simons, *Rethinking Mental States*, 72 B.U. L. REV. 463, 486–90 (1992).



about risk itself. Instead, such accounts deny that reckless defendants must be aware of risk at all, proposing instead that recklessness requires an entirely different kind of psychological attitude, such as culpable indifference.<sup>12</sup> By contrast, though I will deny that recklessness requires defendants to possess some belief directed towards the existence of risk, my account will not abandon the standard approach, found in the Model Penal Code, of understanding recklessness in terms of conscious awareness. In my view, reckless defendants must be aware of risk, but I will advance an alternative account of what mental states are required for the awareness of risk and of when those mental states make a defendant reckless. On that account, a defendant's awareness of risk depends not on beliefs about risk itself but rather on beliefs about the features of her action in virtue of which it is risky.<sup>13</sup> And an individual acts recklessly if performing an action with the features she recognizes her action to have would violate the standard of conduct that a law-abiding person would observe.

Broadly speaking, I will analyze recklessness by considering how the definitions of criminal offenses affect how criminal prohibitions guide individuals' conduct. Criminal offenses defined in terms of awareness provide individuals with a rule to employ when evaluating actions in deliberation: if it is a crime to perform an action believed to have certain features, individuals must not perform actions they believe to have those features. Interpretations of recklessness as intentional risk-creation understand awareness of risk—and thus culpability for recklessness—to require a belief about the degree of risk that an action creates. And since, on this approach, whether an action would be a crime depends on an individual's beliefs concerning the risk it creates, such definitions guide individuals to comply with

---

<sup>12</sup> See, e.g., DUFF, *supra* note 3, at 162 (arguing “that an agent's indifference to a risk which she creates is a matter not of her occurrent feelings, but of the meaning of her actions; that such indifference can be a proper basis for criminal liability; and that it can be displayed by one who is unaware of the risk which she creates is a matter not of her occurrent feelings, but of the meaning of her actions; that such indifference can be a proper basis for criminal liability; and that it can be displayed by one who is unaware of the risk which she creates”); TADROS, *supra* note 11, at 254 (discussing a defendant whose “beliefs are not true, and in forming those false beliefs, [who] has shown himself to be a worthy target of criminal responsibility”); Simons, *supra* note 11, at 482–86 (contrasting belief-based, desire-based, and conduct-based conceptions of recklessness).

<sup>13</sup> This distinction between two kinds of awareness of risk parallels a distinction between two kinds of awareness of, and thus motivation by, moral reasons—namely, the distinction between being motivated by the belief that an action is right or wrong, which philosophers term motivation by morality *de dicto*, and being motivated by the belief that an action has certain features that as a matter of fact make it right or wrong, which philosophers term motivation by morality *de re*. For an overview of the distinction, see James Dreier, *Dispositions and Fetishes: Externalist Models of Moral Motivation*, 61 PHIL. & PHENOMENOLOGICAL RESEARCH 619, 621–22 (2000).

criminal prohibitions by reasoning explicitly about the risks their actions create. As one leading statement of this view puts it, individuals must “choose only those acts for which the risks to others’ interests – as you estimate those risks – are sufficiently low to be outweighed by the interests, to yourself and others, that you are attempting to advance (discounted by the probability of advancing those interests).”<sup>14</sup>

In Part II of this chapter, I will argue against this approach to guiding individuals’ conduct. Obviously, all actions create risk to some degree, and therefore the permissibility of creating a particular risk must depend both on its degree of risk and on the value of the interests it implicates. If individuals must evaluate actions by explicitly reasoning about the risks they create, then, they must be able to estimate how risky different actions are and to determine which actions are permissible in light of those estimates. But it is hopeless to expect individuals actually to employ this reasoning in the ordinary decision-making contexts regulated by criminal law. Outside of highly artificial situations—games of chance, say—individuals lack access to the evidence that would be required to reliably estimate risk. And though individuals have some sense of which outcomes are good or bad, nobody knows the ratios of the values of those outcomes precisely enough to feasibly determine whether an action’s risk of causing a particular harm is outweighed by its benefits. We do not—we could not—actually evaluate actions by first evaluating their risk. Instead, we evaluate actions in deliberation based directly on the features we believe those actions would have. Rules that effectively guide individuals not to create unjustified risks, then, must identify prohibited actions in terms of their actual features.

Part III proposes alternative rules of conduct that specify prohibited actions in terms of their actual features. In my view, recklessness should be defined to promulgate these rules for individuals to employ in deliberation. Using the distinction between risk and uncertainty introduced by John Maynard Keynes and Frank Knight,<sup>15</sup> I will first reframe the decision problems that prohibitions on recklessness ordinarily guide individuals in facing. Interpretations of recklessness as intentional risk-creation understand individuals as facing a risk—that is, a knowable likelihood

---

<sup>14</sup> ALEXANDER, FERZAN, & MORSE, *supra* note 8, at 263.

<sup>15</sup> JOHN MAYNARD KEYNES, A TREATISE ON PROBABILITY 20–43 (1921); FRANK H. KNIGHT, RISK, UNCERTAINTY, AND PROFIT 197–232 (1921).

of causing a particular harm—and consequently promulgate rules that identify how individuals should act given that likelihood. By contrast, in my view individuals ordinarily face uncertainty: they know neither whether their actions will actually cause harm nor the likelihood that those actions will cause harm. Individuals facing uncertainty must employ a different kind of rule to make decisions: without knowledge of how likely their actions are to cause harm, they cannot decide what to do by applying a rule defined directly in terms of that likelihood. What rule, then, should they employ? Because such individuals cannot themselves reliably evaluate the harmfulness of possible actions, they should instead defer to the judgments of others: individuals facing uncertainty should avoid actions that nobody else would perform. In particular, societies informally develop and enforce norms that govern how to perform dangerous activities safely, identifying actions that are excessively risky. Rather than requiring individuals to reason about risk themselves, the prohibitions individuals should follow require them simply to comply with those norms. Or, to employ the language of the unjustly neglected second sentence in the Model Penal Code’s definition, recklessness prohibits “a gross deviation from the standard of conduct that a law-abiding person would observe in the actor’s situation.”<sup>16</sup>

Having identified the rule that prohibitions on recklessness should guide individuals to follow, in Part IV I analyze what mental states an individual must possess to violate that rule in deliberation. Criminal prohibitions enforce rules by defining violations of those rules as offenses; thus, recklessness must be defined so that defendants commit an offense by failing to deliberate in accordance with the relevant social norms. Social norms identify excessively risky actions by describing the features of the actions that individuals are prohibited from performing. To apply a norm in deliberation, then, individuals must determine whether a possible action possesses the features that norm employs to describe prohibited conduct. And they violate the norm by recognizing their action to possess those features, then performing it. Such defendants are reckless: recklessness, then, requires a defendant to recognize his action to possess the features in virtue of which it violated some

---

<sup>16</sup> MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962).

social norm governing dangerous activities. But no further mental state is required. In particular, since prohibitions on recklessness require individuals to defer to social norms rather than reasoning about risk themselves, defendants are not exculpated either by failing to recognize that their action would create risk or by believing that it would not create risk, so long as they do recognize the features in virtue of which it violates a social norm.

### **I. Reasoning About Risk**

One central function of criminal law is to provide legally binding rules of conduct that guide individuals' behavior.<sup>17</sup> It promulgates such rules in part through its definitions of criminal offenses. In particular, offense definitions guide conduct through their descriptions of which actions constitute offenses, since individuals deliberating over whether to perform an action must determine whether it possesses the features identified in those descriptions in order to determine whether it may be performed. Because one consideration relevant to whether an action should be performed is the risk that it will cause an invasion of some legally protected interest, the law should guide individuals to consider risk in evaluating actions. A rule specifying how risk should influence deliberation faces an obvious challenge that other sorts of rules may avoid, however. Risk is ubiquitous: in an uncertain world, there is always some chance that an action will cause harm. Consequently, the law cannot simply instruct individuals not to create risk, which in practice would require individuals to avoid all actions entirely; instead, it must identify how individuals may and may not create it. Thus, the rule of conduct criminal law promulgates to govern risk-taking must distinguish among different risks, identifying when individuals must avoid an action because of the risk it creates and when they are permitted to act despite that risk.

Recklessness is among the law's doctrinal mechanisms for distinguishing permitted from prohibited risk-creation.<sup>18</sup> In the Model Penal Code's formulation, a

---

<sup>17</sup> Arguably, this function is essential to law in general. See Joseph Raz, *The Rule of Law and its Virtue*, in *THE AUTHORITY OF LAW: ESSAYS ON LAW AND MORALITY* 210, 226 (1979) ("The law to be law must be capable of guiding behaviour . . .").

<sup>18</sup> Many criminal offenses may be performed simply by recklessly causing a certain result. *E.g.*, MODEL PENAL CODE § 211.1(1) (AM. LAW INST. 1962) ("A person is guilty of assault if he (a) . . . recklessly causes bodily injury to another."). Besides the definition of recklessness, these definitions have no mechanism for distinguishing

reckless defendant “consciously disregards a substantial and unjustifiable risk.”<sup>19</sup> An offense defined using the mens rea of recklessness, then, identifies a type of risky action that individuals are prohibited from performing, in part by describing the mental state with which those actions are performed: an action is prohibited if the agent acts with a particular mental state—namely, conscious disregard of a substantial and unjustifiable risk—and her action satisfies the other components of the offense.<sup>20</sup> The definition of recklessness thereby identifies a rule for individuals to employ in deliberation when evaluating potential actions: whether someone may perform a particular action depends on whether, in light of her mental attitudes towards the risks that action creates, she would act with that mental state, and therefore be reckless, were she to perform it.

If criminal prohibitions framed in terms of recklessness employ the agent’s own mental states to distinguish between permissible and prohibited acts of risk-creation, then the kind of mental state required for recklessness determines how the law will guide individuals’ behavior. But the Model Penal Code’s definition is strikingly unhelpful in identifying that mental state. The explicit definition it gives of acting recklessly merely states that a defendant “acts recklessly . . . when he consciously disregards a substantial and unjustifiable risk.”<sup>21</sup> The phrase “conscious disregard” does little to identify which mental state is required for recklessness, though, beyond specifying that it must be among the conscious rather than unconscious contents of a defendant’s mind.<sup>22</sup> And while the Model Penal Code’s commentaries do explain that a defendant consciously disregards a risk by being

---

permitted from prohibited risks; thus, the definition of recklessness must on its own determine whether the action that caused injury was sufficiently risky to constitute a crime. Other offenses, however, do include elements relating to risk besides the mens rea requirement—for example, the offense of “Recklessly Endangering Another Person” requires a defendant to “recklessly engage[] in conduct which places or may place another person in danger of death or serious bodily injury.” *Id.* § 211.2. It is not obvious how the definition of recklessness interacts with the additional requirement that the defendant’s conduct must be dangerous (or perhaps only possibly dangerous, if the conduct only “may place” another in danger). My arguments in this chapter will focus on the definition of recklessness, though; I will not analyze other doctrines that concern related concepts like danger.

<sup>19</sup> MODEL PENAL CODE § 2.02(2)(c).

<sup>20</sup> Of course, because many offenses of recklessness are defined to require that harm occur, a reckless action that luckily causes no harm may not be criminally punishable. Nonetheless, such definitions do not promulgate a rule permitting harmless, reckless actions: since individuals cannot distinguish between equally reckless actions that will and will not in fact produce harm, a prohibition on the former requires individuals to avoid the latter, as well.

<sup>21</sup> MODEL PENAL CODE § 2.02(2)(c).

<sup>22</sup> Kenneth W. Simons, *Should the Model Penal Code’s Mens Rea Provisions Be Amended?*, 1 OHIO ST. J. CRIM. L. 179, 192–95 (2003) (considering conceptual difficulties concerning the notion of “conscious” disregard).

aware of it,<sup>23</sup> that clarification is hardly an improvement, for it is uncertain both what kind of mental state awareness is and what exactly an individual must be aware of for that awareness to be awareness of risk.<sup>24</sup>

Clarifying the nature of recklessness, then, requires clarifying the nature of awareness of risk. Since awareness of risk is a particular kind of mental state that somehow represents the risk an action creates, understanding it might be approached primarily as a problem in psychology and epistemology, requiring analysis of the psychology of awareness and the epistemology of risk.<sup>25</sup> But the connection I have noted between the awareness of risk and the action-guiding function of criminal law suggests that awareness of risk might instead be approached as a normative problem: the mental state required for a defendant to be guilty of recklessness might be identified by considering how individuals in fact ought to reason about risk in deciding what to do. The law's definition of recklessness distinguishes permitted and prohibited forms of risk-creation: whether an agent would be reckless in acting determines whether she may act. Consequently, whatever mental state defendants must possess to be found reckless will play an important role in guiding individuals' conduct: because the law forbids risk-creation when it is reckless, it requires individuals not to act just in case they possess whatever mental state is required for recklessness. And if the mental state required for recklessness specifies when individuals must avoid actions because of their risk, an account of which rules individuals should follow to determine whether an action is too risky to perform will

---

<sup>23</sup> MODEL PENAL CODE § 2.02 cmt. 3 (AM. LAW INST. 1962) (“[Recklessness] resembles acting knowingly in that a state of awareness is involved . . .”). Some legal systems specify that recklessness requires awareness in the operative legal rules themselves, not in their commentaries: under English law, for example, “A person acts recklessly within the meaning of section 1 of the Criminal Damage Act 1971 with respect to (i) a circumstance when he is aware of a risk that it exists or will exist; (ii) a result when he is aware of a risk that it will occur . . . .” *R v. G and another*, [2003] UKHL 50, [2004] 1 AC (HL) 1034, [41] (Lord Bingham of Cornhill) (appeal taken from Eng.).

<sup>24</sup> *E.g.*, Husak, *supra* note 7, at 207 (describing “the nature of awareness” as a “pivotal topic [that] is radically under-theorized”); Kenneth W. Simons, *Culpability and Retributive Theory: The Problem of Criminal Negligence*, 5 J. CONTEMP. LEGAL ISSUES 365, 384 (1994) (“[T]he minimum requirements of ‘awareness’ . . . are very difficult to identify.”).

<sup>25</sup> *E.g.*, STARK, *supra* note 6, at 90–122 (analyzing awareness in terms of a dispositional account of belief); R.A. Duff, *Caldwell and Lawrence: The Retreat from Subjectivism*, 3 OXFORD J. LEGAL STUD. 77, 80 (1983) (analyzing awareness of risk by positing a distinction between “explicit,” “tacit,” and “latent” knowledge); Ferzan, *supra* note 8, at 635–41 (arguing that recklessness requires “introspective awareness” rather than “behavioral awareness”); Husak, *supra* note 7, at 207–210 (analyzing awareness in terms of a dispositional account of belief); Sarch, *supra* note 10, at 727–30 (analyzing awareness in terms of knowledge); Williams, *supra* note 8, at 84 (“[A] person is unaware of the risk only if he believes that his act is safe.”).

itself identify what mental state the law should require defendants to possess in order to be guilty of recklessness.

In general, we decide how to act by first gathering and processing information about our circumstances and our available actions, then by evaluating those actions and choosing one to perform. Criminal law promulgates rules that govern this process by identifying actions that are prohibited and thus may not be chosen. Its rules concerning risk, in particular, must state some criterion for individuals to apply in deliberation to identify actions that may not be chosen because of their risk: the law requires individuals to reason about risk by applying that criterion to possible actions and avoiding the actions that it excludes. To comply with those rules, individuals deliberating over an action must determine whether it meets the conditions the criterion uses to identify prohibited risk-taking. And an agent breaches a rule governing how to reason about risk if her action is inconsistent with that rule—if, that is, she would not have performed the action had she applied the rule in deliberation and complied with its conclusions.<sup>26</sup>

The conclusions an individual reaches by applying a rule in deliberation on a particular occasion, though, depend on her other mental states. Roughly speaking, any rule agents employ in reasoning allows them to draw certain conclusions from certain premises. The results of applying the rule, then, depend on whether the agent accepts its premises, for only if so will the rule instruct her to accept its conclusions. Rules governing the creation of risk are no different: because an agent can determine whether an action satisfies the rule's criterion for prohibited risk-taking only using her information about that action, the results of applying the rule on a particular occasion will depend on the information available to her. Any rule guides the choices

---

<sup>26</sup> This argument takes an individual to violate a rule by performing an action that would have been excluded had the rule been applied in deliberation. But another possibility exists: perhaps an individual violates a rule not based on how it would have influenced deliberation, had it been applied, but rather only if he actually does apply the rule in deliberation but disregards its conclusion by choosing an action it excludes. On this approach, individuals violate a rule only if they recognize, in acting, that their action would violate it; if they did not actually apply the rule in their deliberation, it is not enough that the rule would have forbidden that action had it been applied. Some philosophers have defended an account of moral blameworthiness along these lines. *See, e.g.*, Gideon Rosen, *Culpability and Ignorance*, 103 *PROC. ARISTOTELIAN SOC'Y* 61 (2003); Gideon Rosen, *Skepticism about Moral Responsibility*, 18 *PHIL. PERSP.* 295 (2004); Michael J. Zimmerman, *Moral Responsibility and Ignorance*, 107 *ETHICS* 410 (1997). But since criminal law does not generally require individuals to recognize that their action violates a rule in order to be culpable for violating it, *see* MODEL PENAL CODE § 2.02(9) (AM. LAW INST. 1962), I will not consider an account of recklessness based on this notion of when a rule is violated.

reached in deliberation based on the information employed in deliberation: it takes as an input certain of the agent's beliefs and other representational mental states, then yields a conclusion as to whether an action may be performed. Thus, whether an action violates the rules governing reasoning about risk depends on the mental states the agent employed in reasoning about risk. She violates the rule only by performing an action she would have identified as prohibited had she applied the rule in deliberation—that is, only if applying the rule to the mental states she actually possessed would entail that acting would be prohibited. If she lacked those mental states, her risky actions would not violate the rule: she would apply the rule correctly to her inputs, and her action would be flawed only because those inputs themselves were in some way mistaken.<sup>27</sup>

An individual violates the rules governing reasoning about risk, then, if applying those rules given her mental states would yield the conclusion that her action was prohibited. Criminal prohibitions, of course, enforce rules by defining violations to be offenses. Thus, if the definition of recklessness is to promulgate rules identifying when risk-taking is prohibited, it should be defined so that agents who violate the rules governing reasoning about risk are reckless—in particular, so that individuals act recklessly if they would have identified their action as forbidden had they applied those rules in deliberation. Thus, those rules can identify what mental state is required for recklessness—namely, individuals are reckless if they possess mental states from which the rule would require them to conclude that an action is too risky to perform. And because the content of those rules determines the identity of those mental states, a theory of recklessness depends on an account of how individuals should reason about risk in deliberating over what to do. Individuals are reckless if, given their mental states, appropriate reasoning about risk would identify their action as excessively risky. My approach in this chapter, then, will be to construct an account of recklessness by considering the normative question of what

---

<sup>27</sup> The use of those inputs may itself violate some other rule governing deliberation. For example, in chapter 3 of this dissertation I argue that negligence consists in the failure to engage in adequate inquiry before acting; obviously, then, an individual's negligence may affect the information available to her in applying other rules in deliberation.



sort of reasoning about risk the law should require individuals to employ in evaluating actions.

Accounts of recklessness as intentional risk-creation can be understood as an application of this approach. First, they propose a rule for individuals to employ in deliberation that distinguishes permitted from prohibited risks based on their degree or magnitude. This suggestion follows quite naturally from definitions of recklessness: the Model Penal Code, for example, holds defendants to be reckless if they consciously disregard a “substantial and unjustifiable” risk.<sup>28</sup> Controversies exist over how substantiality and unjustifiability should each be understood, as well as over whether they are distinct requirements, but whether a risk is substantial and unjustifiable plausibly depends in some way on how risky it is.<sup>29</sup> This rule, then, states how individuals must reason about risk in deliberating over whether to perform an action: they must evaluate actions by identifying the magnitude of the risk they would create. In particular, if the risk that an action creates is substantial and unjustifiable in its magnitude, the rule judges that action impermissible, and individuals must refrain from performing it.

This account of the rule individuals must employ in reasoning about risk, in turn, identifies what mental state is required for an action to violate the rule. Individuals violate a rule by failing to act as it demands—that is, by performing an action that it would have identified as prohibited if it had been applied in deliberation. And because the rule proposed to govern reasoning about risk identifies prohibited risks based on their magnitude, the key input that determines what conclusions would follow from applying that rule is the magnitude ascribed to a risk: an agent would identify an action as prohibited by applying the rule just in case the degree of risk he believed it would create exceeds the threshold the rule employs to identify prohibited risks. Note that individuals need not further believe that a risk of that degree would be prohibited; on this proposal rule violations do not require an individual to have any belief about what conclusion would result from applying the rule. Instead, rule

---

<sup>28</sup> MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962). Similar terminology is employed in other modern legal systems. *See, e.g.*, R v. G and another, [2003] UKHL 50, [2004] 1 AC (HL) 1034, [41] (Lord Bingham of Cornhill) (appeal taken from Eng.) (“A person acts recklessly . . . when . . . it is unreasonable to take the risk.”).

<sup>29</sup> For discussion of these questions, see Treiman, *supra* note 4, at 334–39.

violations occur when the rule would judge acting to be prohibited, were it applied; to violate it individuals need not actually apply it and recognize an action to be prohibited.<sup>30</sup> So long as the degree of risk an agent expects an action to create exceeds the threshold specified in the rule, she would have identified the action as prohibited by applying the rule in deliberation. Consequently, the definition of recklessness must require agents to possess some mental state—a belief, say, or an intention—directed towards the degree of risk that an action creates, and that risk must in fact exceed the threshold of permissibility.<sup>31</sup> Defining recklessness in this way identifies the rule individuals must employ in deliberation: performing an action is prohibited if one expects it to create a degree of risk that exceeds a certain threshold, so individuals may not select any action in deliberation if they expect it to create a degree of risk that exceeds that threshold.<sup>32</sup>

This account of how the prohibition on reckless conduct requires individuals to incorporate considerations of risk into their deliberations about what to do, furthermore, neatly extends the paradigm established by prohibitions on intentional wrongdoing to cover offenses defined in terms of risk. Criminal offenses defined using kinds of mens rea like intention, knowledge or purpose guide individuals' deliberation by promulgating rules that exclude actions based on features like their results or circumstances.<sup>33</sup> Such prohibitions influence conduct by requiring individuals not to perform actions when they possess mental states directed towards

---

<sup>30</sup> See *supra* note 26. Theorists who defend widely different accounts of recklessness have generally agreed that reckless defendants need not believe that the risk their action would create would be unjustified; instead, it suffices for them to correctly recognize the degree of risk their action would create, so long as a risk of that degree is unjustified. See Treiman, *supra* note 4, at 365–68.

<sup>31</sup> For some of the many proposals to understand recklessness along these lines, see ALEXANDER, FERZAN, & MORSE, *supra* note 8, at 28 (“[A]n actor is reckless if he believes he is imposing a level of risk that would be unjustifiable . . . .”); STARK, *supra* note 6, at 122 (“The reckless defendant . . . merit[s] the ascription to him of a belief that there is a specific, substantial risk attendant upon  $\phi$ -ing . . . .”); Ferzan, *supra* note 8, at 598 (characterizing as reckless those who “engage in risky behavior, fully acknowledging the possibility that their actions could cause substantial and unjustifiable harm to others”); Claire Finkelstein, *Responsibility for Unintended Consequences*, 2 OHIO ST. J. CRIM. L. 579, 595 (2005) (“[T]he agent who consciously disregards a substantial risk is responsible for the consequence that eventuates from that risk only if he was aware not only of the risk’s existence, but also of its substantiality.”); Husak, *supra* note 7, at 208 (“[T]he reckless defendant must be aware of a risk, aware that it is substantial, and also aware of the facts that make it unjustifiable . . . .”); Sarch, *supra* note 10, at 729 (proposing that “recklessness requires the true belief that one imposes [substantial] risks”); Williams, *supra* note 8, at 90–91 (“If the defendant knows that he is running a given risk, and if that risk is objectively adjudged to be unjustifiable, then the defendant is guilty of recklessness . . . .”).

<sup>32</sup> ALEXANDER, FERZAN, & MORSE, *supra* note 8, at 263 (“[C]hoose only those acts for which the risks to others’ interests – as you estimate those risks – are sufficiently low to be outweighed by the interests, to yourself and others, that you are attempting to advance (discounted by the probability of advancing those interests).”).

<sup>33</sup> See, e.g., MODEL PENAL CODE § 2.02(2)(a)–(b) (AM. LAW INST. 1962).

certain features of those actions. Thus, an individual violates the rule, and thereby performs an offense, by acting with that mental state—that is, with an intention, belief, or purpose that the action performed possesses a feature in virtue of which it would be prohibited. Interpretations of recklessness as intentional-risk creation promulgate the same sort of rule for agents to apply in deliberation—one that excludes an action when agents possess a mental state directed towards a specific feature of that action, and that is violated if they act with such a mental state. The only difference, of course, is what specific feature of actions is involved: while other kinds of offenses involve mental states directed towards some state of affairs in the world, like an action's results or circumstances, offenses of recklessness involve a mental state directed towards the risk it creates. But the underlying model of how individuals should deliberate about what to do is the same: individuals must decide whether to perform an action based on their own assessment of whether it possesses some specific feature in virtue of which it would be prohibited.

Given the importance of the paradigm of intentional wrongdoing to criminal law, the appeal of extending it to cover risk-creation is understandable. I will argue, however, that criminal law cannot effectively regulate actions based directly on their risk. Defining offenses in terms of a mental state whose object is a specific feature of the action to be performed identifies a particular kind of rule individuals must employ in deliberating over what to do—one that evaluates possible actions based on whether the agent intends or believes them to have that feature. Such offenses thereby guide individuals to deliberate by evaluating different actions in light of the features they are expected to have. In general, this strategy succeeds when used to define offenses with the mens rea of intention, purpose, or knowledge. Individuals are reasonably effective at identifying actual features of their actions, like their results or circumstances; thus, the actions individuals believe or intend to have a particular harmful feature likely will have that feature. Furthermore, rules that prohibit actions based on those features are straightforward and easy to apply, since they ordinarily simply require individuals to avoid an action if it has certain features. Consequently, the law can guide conduct effectively by defining offenses using mental states directed towards harmful features of actions: such prohibitions promulgate rules that

individuals can effectively and accurately apply in deliberation, and thus actions individuals identify as prohibited by applying those rules are, in fact, actions that likely will be harmful.

Requiring individuals to deliberate over actions based on mental states directed towards those actions' actual features is effective, then, because individuals can reason effectively about the features actions are expected to have. But a key difference exists between risk and the features of actions ordinarily used to define offenses. In general, those features apply to actions with a binary logic: they either do or do not characterize particular actions. Acting either will or will not cause injury, for example, just as it either will or will obtain the property of another by deception. Thus, a mental state represents any action either as having that feature or not. By contrast, risks come in degrees, and mental states may represent them as taking on a wide range of values. Consequently, reasoning with mental states directed towards risk is a much more challenging task. First, since risks come in degrees, individuals must identify the magnitude of particular risks in order to reason about them. But it is very challenging to accurately identify how risky a particular action is. Thus, while reasoning about the actual features of actions generally employs accurate inputs, the inputs to reasoning about risk will involve inaccurate judgments about the magnitudes of risks. Second, because the magnitude of a risk determines how it should influence behavior, reasoning about risk requires rules complex enough to allow risks of different magnitudes to have different influences on choice. And rules complex enough to require different decisions when risks differ in degree are much harder to apply than rules that simply instruct individuals not to perform actions with particular features.<sup>34</sup>

Unlike explicit reasoning about the actual features of actions, in short, explicit reasoning about risk is hard to do well. And because it is hard to do well, individuals are unlikely to do it well: in particular, they are unlikely to accurately identify which

---

<sup>34</sup> To be sure, many factors that come in degrees, besides risk, are normatively relevant to decision-making. The permissibility of driving on the highway, for example, depends on one's speed, which comes in degrees; the permissibility of using force in self-defense depends on the threat that one faces, which comes in degrees. My arguments here for why criminal law cannot easily regulate actions by their risk can be extended to all cases, such as these, in which a normatively relevant property of actions is graded rather than binary. But I will not attempt to identify the principles governing whether the solution I propose in this chapter to analyze recklessness could be extended to these other graded properties, as well.

actions are in fact risky enough that they should be avoided. Thus, while criminal law can effectively guide individuals to avoid harm by promulgating rules for deliberation that evaluate actions based on their actual features, it cannot effectively guide conduct by promulgating rules that evaluate actions based on the risk they create, for individuals are unlikely to apply those rules accurately in deliberation. As I have argued, though, definitions of recklessness that understand conscious awareness in terms of an explicit mental state directed towards risk guide individuals to reason about risk in precisely this manner: by prohibiting risk-creation based on the magnitude of the risk that defendants take themselves to be creating, such definitions guide individuals to deliberate over actions by first identifying the magnitude of the risk they will create, then evaluating them on that basis. If such evaluations are likely to be inaccurate, criminal prohibitions should not guide individuals to employ them. Instead, an account of recklessness requires a better theory of how individuals should reason about risk: in particular, it should be defined so that defendants act recklessly by violating the rules they actually should apply in deliberation to avoid excessively risky actions.

#### A. *Quantitative Reasoning About Risk*

On my account, the challenges that individuals face in reasoning about risk arise from the fact that risk—unlike actual properties of actions—comes in degrees. But there are different ways for an agent to measure and represent the degree of a risk when deliberating about it. And since the degree of a risk affects the influence it rationally ought to exert on individuals' choices, different methods of measuring and representing risk will require different methods of reasoning about it. One general approach to measuring and representing risk is to quantify it. An individual would be aware of risk, on this approach, only by possessing some quantitative estimate of it. In general, understandings of recklessness as intentional risk-creation require individuals to evaluate actions by determining whether the risk those actions will create is substantial and unjustifiable in magnitude. On a quantitative approach to representing risk, then, that sort of evaluation would be quantitative: a risk would be a substantial and unjustifiable if it exceeded some quantitative threshold. And individuals would follow a distinctive procedure in reasoning about risk, first

quantifying the magnitude of the risk an action would create, then applying some rule to determine whether a risk of that magnitude would be substantial and unjustifiable. In this section, I will argue that individuals cannot reason about risk effectively in this manner. Neither of the two steps involved in this sort of evaluation of risk would ordinarily be performed accurately: individuals cannot accurately quantify the degree of the risk an action creates, and they cannot in general accurately apply rules that evaluate actions based on quantified estimates of those risks. Even if some quantitative threshold plausibly does distinguish between risks individuals should and should not create, individuals are unlikely to accurately identify which actions cross that threshold, in practice. And because this method of reasoning about risk will in general be inaccurate, criminal prohibitions should not guide individuals to employ it.

What sort of quantitative threshold might the law employ to identify which risks are substantial and unjustifiable and thus reckless to create? One common suggestion is to employ expected-utility models familiar from decision theory.<sup>35</sup> These approaches model rational choice by assigning a value to each action available in a particular decision—its expected utility—based on the effects it is expected to produce in the world. An action’s expected utility depends jointly on the worth and the likelihood of its effects: each possible outcome of each possible action must be assigned both a quantitative utility, which captures its value, and a quantitative probability, which captures the probability that the action will produce it. The expected utility of the action, then, is the sum of the values of all of its possible effects, each multiplied by its likelihood of actually occurring. Agents must choose the action with the highest value, thus defined.<sup>36</sup>

Expected-utility reasoning is most often understood as an account of rational choice, formally characterizing what it is for an agent to choose the action that is likeliest to advance his own preferences given his own beliefs about his options and his circumstances. Obviously, though, prohibitions on recklessness do not aim merely

---

<sup>35</sup> For suggestions that the unjustifiability of risks should be understood along these lines, see ALEXANDER, FERZAN, & MORSE, *supra* note 8, at 263; STARK, *supra* note 6, at 21–22; and Eric A. Johnson, *Beyond Belief: Rethinking the Role of Belief in the Assessment of Culpability*, 3 OHIO ST. J. CRIM. L. 503, 507 (2006).

<sup>36</sup> For an overview of this account of rational choice, see generally R.A. Briggs, *Normative Theories of Rational Choice: Expected Utility*, STAN. ENCYCLOPEDIA PHIL. (Edward N. Zalta ed., Fall 2019 ed.), <https://plato.stanford.edu/archives/fall2019/entries/rationality-normative-utility/>.

to prohibit irrationality; rather, like the rest of criminal law they require individuals to avoid actions that invade or risk invading certain important interests. But expected-utility models of reasoning can be adapted to serve this function. Rational agents evaluate each outcome according to the utility that they themselves assign it in order to choose the action that best achieves their own preferences; by contrast, since criminal law aims to avoid invasions of legally protected interests, it must require individuals to treat invasions of those interests as undesirable. Rather than subjectively evaluating those outcomes, then, agents must assign appropriately low utilities to such invasions in computing expected utilities. Furthermore, it is implausible that prohibitions on recklessness should require individuals always to maximize expected utility, since criminal law ordinarily does not prohibit all actions that are in any way suboptimal but rather prohibits only those that are seriously wrongful. Thus, only actions with sufficiently low expected utility should be defined to be reckless. It is not entirely clear how that threshold should be set, but presumably actions must be expected to produce more harm than good to be prohibited, perhaps much more.

An action would create a substantial and unjustifiable risk, then, if the magnitude of that risk is great enough that the harm the action is expected to produce outweighs its expected benefits. This interpretation of the threshold at which creating a risk would be reckless identifies a rule governing how individuals should deliberate about risk: they may not perform actions with sufficiently low expected utility, calculated using the utilities that the law assigns to each outcome. In Larry Alexander, Kimberly Kessler Ferzan, and Stephen Morse's description, such rules instruct individuals to "choose only those acts for which the risks to others' interests – as you estimate those risks – are sufficiently low to be outweighed by the interests, to yourself and others, that you are attempting to advance (discounted by the probability of advancing those interests)."<sup>37</sup> Applying this rule requires individuals to assign probabilities to the chance that an action will produce each possible outcome, since those probabilities set the discount rate to be employed in computing expected utility. If the expected utility of an action would be sufficiently low, given the utilities that

---

<sup>37</sup> ALEXANDER, FERZAN, & MORSE, *supra* note 8, at 263.

the law assigns to outcomes and the probabilities the agent assigns to them, applying this rule would identify the action as prohibited. And because a defendant is reckless if he acts despite possessing mental states from which he could have inferred, using the rule, that an action is prohibited, he acts recklessly if his assignment of quantified probabilities to outcomes entails that the expected utility of acting would be sufficiently negative.

This approach clearly does involve a rule that identifies the risks individuals may not create. And it is easy to understand why that rule might be appealing: surely risks whose costs sufficiently outweigh their benefits should be prohibited. But if an action's expected utility determines whether it is excessively risky, individuals must reason about risk in a particular manner when deliberating over whether to perform an action—namely, by computing its expected utility, since that determines whether it may be performed. In the vast majority of real-life situations, however, it would be “utterly ridiculous” and “preposterous” to apply explicit expected-utility reasoning, in the words of Leonard Savage, one of the founders of decision theory.<sup>38</sup> While in theory individuals should avoid actions whose harms, discounted by their likelihood of occurring, outweigh their benefits, similarly discounted, in practice individuals have no reliable way to accurately identify which actions those are. First, they lack any reliable method of accurately identifying either the appropriate values to assign to those harms and benefits or the appropriate probabilities by which to discount them. Second, the complex, mathematical process of computing expected values is taxing enough that real, cognitively limited human beings cannot plausibly perform it when making actual decisions. Expected-utility reasoning is not an effective way for ordinary individuals deliberating over what to do to identify excessively risky actions. But individuals must employ reasoning of that sort in order to avoid performing an offense if recklessness is defined so that whether an act constitutes an offense depends on its expected utility given the agent's quantitative estimates of its probability of causing harm. Consequently, the law should reject such definitions of recklessness.

---

<sup>38</sup> LEONARD J. SAVAGE, THE FOUNDATIONS OF STATISTICS 16 (1954).



First, expected-utility reasoning is ill-suited to ordinary deliberation because individuals ordinarily lack the mental states required to compute the expected utilities of their actions. The expected utility of an action depends closely on how likely it is to cause harmful or beneficial outcomes. A two percent probability of causing harm, for example, reduces an action's expected value only one fourth as much as an eight percent probability of causing harm, a difference that might affect the permissibility of a wide range of possible actions. Because the precise probability assigned to an outcome determines its expected utility, and thus whether individuals may perform it, individuals can reason about risk using expected utilities only given quantified estimates of the probability that an action will cause harm. It is wildly unrealistic, however, to suggest that agents typically possess such probability estimates. In most of the ordinary decisions we face in daily life we do not first quantify probabilities before making a choice: we do not decide whether to take an umbrella by first quantifying the probability of rain, nor do we decide whether to make a particular turn while driving by first quantifying the probability of a crash.<sup>39</sup> Explicit, quantified probability estimates are not available in ordinary reasoning about risk. And without them, we simply cannot employ expected-utility reasoning in deliberation, since an essential step in that approach—discounting the value of harms and benefits outcomes by their probability—cannot be performed without quantitative estimates of probability.

Of course, requiring individuals to deliberate using rules that employ quantitative probabilities might still be reasonable if individuals could quantify probabilities but merely choose not to: the law might require individuals to reason about risk using expected utilities by requiring them to quantify the probabilities that acting will have various outcomes. But individuals do not ordinarily quantify

---

<sup>39</sup> *E.g.*, KEYNES, *supra* note 15, at 22–32 (giving examples of decisions in which the existence of precise probability estimates is highly implausible); KNIGHT, *supra* note 15, at 210 (“The ordinary decisions of life are made on the basis of ‘estimates’ of a crude and superficial character. In general the future situation in relation to which we act depends upon the behavior of an indefinitely large number of objects, and is influenced by so many factors that no real effort is made to take account of them all, much less estimate and summate their separate significances. It is only in very special and crucial cases that anything like a mathematical [exhaustive and quantitative] study can be made.”); PAUL WEIRICH, REALISTIC DECISION THEORY: RULES FOR NONIDEAL AGENTS IN NONIDEAL CIRCUMSTANCES 60 (2004) (“Nonquantitative cases arise when the grounds of quantitative probability and utility assignments are absent. This happens commonly . . . .”); James M. Joyce, *A Defense of Imprecise Credences in Inference and Decision Making*, 24 PHIL. PERSP. 281, 282–83 (2010) (“As many commentators have observed, numerically sharp degrees of belief are psychologically unrealistic.” (citation omitted)).

probabilities because, as Keynes and Knight argue, they cannot do so rationally: their evidence does not identify any particular quantitative probability as the likelihood that a particular outcome will occur. In general, quantified probability estimates may be rationally justified by evidence that identifies the proportion of equally probable possible states of affairs in which an outcome occurs.<sup>40</sup> Evidence can identify that proportion in different ways. In some cases, symmetries in the world identify a set of equally likely possibilities that may be analyzed theoretically to determine the proportion in which some outcome occurs. Such symmetries are common in games of chance, for example, justifying the judgments that a fair die is equally likely to land on any of its six sides, or that a fair coin is equally likely to land on either face.<sup>41</sup> Alternatively, if a process occurs sufficiently many times, a statistical analysis of its outcomes may quantify probabilities empirically (at least approximately): if each experimental result is assumed to represent an equally likely possibility, then probabilities may be computed by analyzing the frequencies with which certain outcomes occur among those possibilities. Actuarial science, for example, employs this approach to model life expectancies or to price insurance policies.<sup>42</sup> Each of these kinds of evidence rationally justifies quantifying the likelihood of harm because it identifies the proportion of possibilities in which harm occurs.

As both Keynes and Knight stress, though, it is quite rare that evidence of this sort is available.<sup>43</sup> It is possible to calculate precise probabilities only to the extent that the evidence defines some set of equally probable possibilities and identifies how frequently among them some outcome occurs.<sup>44</sup> And although the evidence certainly

---

<sup>40</sup> KEYNES, *supra* note 15, at 70 (“[A] numerical measure can actually be obtained in those cases only in which a reduction to a set of exclusive and exhaustive equiprobable alternatives is practicable.”); KNIGHT, *supra* note 15, at 220 (“If the chance of any particular result is more or less than one half, it is held to be axiomatic that there is a greater number of possible alternatives which yield this result [or do not yield it] than of the other kind; the alternatives themselves must be equally probable. The whole mathematical theory of probability is obviously a simple application of the principles of permutations and combinations for finding out the number of alternatives.”).

<sup>41</sup> See KNIGHT, *supra* note 15, at 214–15 (giving a fair die as an example of an a priori probability).

<sup>42</sup> KEYNES, *supra* note 15 at 30–31 (mentioning that actuarial science produces quantified probabilities through statistical analysis); KNIGHT, *supra* note 15, at 213 (illustrating the possibility of statistically quantifying risks through the examples of fire loss and of champagne bottles bursting in production).

<sup>43</sup> KEYNES, *supra* note 15, at 182 (“It is evident that the cases in which exact numerical measurement is possible are a very limited class . . . .”); KNIGHT, *supra* note 15, at 226 (arguing that the inability to precisely identify probabilities “obviously applies to the most of conduct”).

<sup>44</sup> KEYNES, *supra* note 15, at 182 ([E]xact numerical measurement is . . . generally dependent on evidence which warrants a judgment of equiprobability by an application of the Principle of Indifference.”).

does sometimes do that, it does so rarely in the decisions we ordinarily make in daily life. Typically, ordinary, everyday decisions are *sui generis*: as Knight explains, “The essential and outstanding fact is that the ‘instance’ in question is so entirely unique that there are no others or not a sufficient number to make it possible to tabulate enough like it to form a basis for any inference of value about any real probability in the case we are interested in.”<sup>45</sup> An individual deciding whether to perform a particular action—whether to bring an umbrella given visible black clouds,<sup>46</sup> say, or whether to expand a particular business enterprise in a particular manner<sup>47</sup>—cannot rationally assign a quantitative probability to a particular outcome: her evidence will not ordinarily define any set of similar decisions made in similar circumstances, and even if it did she would have no realistic way to analyze them to quantify how frequently different outcomes resulted from choosing different options.<sup>48</sup> She cannot rationally assign quantified probabilities to the outcomes of acting, and without them she cannot compute whether the expected utility of creating a particular risk would be negative. Of course, she could still assign quantitative probabilities to outcomes without evidence justifying that assignment. But those probabilities would be random guesses, and thus they would likely be inaccurate: in general, they will bear little relationship to the frequency with which actions actually do produce harm. Actions with negative expected utility, as computed using those probabilities, will not generally produce a greater balance of harms than benefits than other actions, and a rule guiding individuals to evaluate actions in deliberation on that basis will not effectively reduce the harms that their actions cause in the aggregate.

In part because ordinary individuals cannot plausibly assign a quantitative probability to the outcomes of their actions based on the evidence available in ordinary deliberation, some decision theorists have developed a very different

---

<sup>45</sup> KNIGHT, *supra* note 15, at 226; *see also* KEYNES, *supra* note 15, at 30 (“[W]here general statistics are available, the numerical probability which might be derived from them is inapplicable because of the presence of additional knowledge with regard to the present case.”); KNIGHT, *supra* note 15, at 231 (describing decisions in which “there is no possibility of forming in any way groups of instances of sufficient homogeneity to make possible a quantitative determination of true probability. Business decisions, for example, deal with situations which are far too unique, generally speaking, for any sort of statistical tabulation to have any value for guidance. The conception of an objectively measurable probability or chance is simply inapplicable.”).

<sup>46</sup> KEYNES, *supra* note 15, at 31–32.

<sup>47</sup> KNIGHT, *supra* note 15, at 226.

<sup>48</sup> *Id.* at 231 (“[T]here is no possibility of forming in any way groups of instances of sufficient homogeneity to make possible a quantitative determination of true probability.”).

strategy for understanding decisions quantitatively. On this approach, an individual does not reason to a probability estimate from her evidence; rather, her estimate is revealed by her preferences over actions. The probability she assigns to an outcome would thus depend on her dispositions to make various actual and hypothetical choices.<sup>49</sup> Different accounts exist of exactly which dispositions determine those probability estimates. The most straightforward approach associates probability judgments with betting dispositions: since an agent's willingness to bet that a particular outcome will occur rationally ought to depend on how likely she takes that outcome to be, her dispositions to accept different bets can justify attributing a particular probability judgment to her.<sup>50</sup> Someone who would risk up to seventy-five cents in a bet that pays a dollar takes the likelihood of the bet's paying off to be seventy-five percent.<sup>51</sup> More rigorous approaches show that if an agent's preferences across hypothetical pairwise decisions satisfy certain axioms, which essentially require those preferences to be consistent, an assignment of utilities and probabilities to each outcome exists such that her choices maximize expected utility relative to that assignment.<sup>52</sup> She behaves as if she accepted those probability and utility assignments. Since this approach argues that we should understand an individual's beliefs about risk to be those that are revealed by her behavior, not those she produces by reasoning about her evidence, it concludes that an individual accepts whatever probability assignment would explain her actual choices.

Whatever the merits of the revealed-preference approach in general, though, the quantitative probabilities that it attributes to individuals cannot plausibly be employed to evaluate the actions they might perform. Certainly, decision theory provides a simple, powerful, and elegant mathematical formalism for describing individual behavior, reducing the full range of an individual's decision-making

---

<sup>49</sup> For an overview of this approach to decision theory, see generally Jake Chandler, *Descriptive Decision Theory*, THE STANFORD ENCYCLOPEDIA OF PHILOSOPHY (Edward N. Zalta ed., Winter 2017 ed.), <https://plato.stanford.edu/archives/win2017/entries/decision-theory-descriptive/>. Key works in its development include Bruno De Finetti, *La Prévission: Ses Lois Logiques, Ses Sources Subjectives*, 7 ANNALES DE L'INSTITUT HENRI POINCARÉ 1 (1937); Frank Plumpton Ramsey, *Truth and Probability*, in THE FOUNDATIONS OF MATHEMATICS AND OTHER LOGICAL ESSAYS 156 (R.B. Braithwaite ed., 1931); and SAVAGE, *supra* note 38.

<sup>50</sup> De Finetti, *supra* note 49, at 6–7; Ramsey, *supra* note 49, at 172–75.

<sup>51</sup> E.g., De Finetti, *supra* note 49 at 6. Of course, this approach ignores the fact that a real-life agent's willingness to bet may depend on considerations other than his estimates of probabilities—whether he enjoys gambling, say.

<sup>52</sup> See generally SAVAGE, *supra* note 38. For a concise overview of Savage's approach, see Chandler, *supra* note 49, at 1.

dispositions to a single utility function and a single credence function. It is no surprise, then, that it has become important in the behavioral sciences.<sup>53</sup> But criminal law does not aim to describe behavior but rather to provide rules that guide it. If it defines recklessness in terms of the quantified probability that an action will cause harm, then, those probabilities must be accessible to individuals in their reasoning, so that individuals can determine how to act based on the risk that an action would cause harm. And even though the revealed-preference approach does ascribe quantitative probabilities to individuals, individuals cannot ordinarily use those probabilities to guide conduct their effectively. Individuals will rarely be consciously aware of those probabilities, which will in any event will generally be too inaccurate to constitute a reasonable basis for deliberation and choice.

In deliberation, agents consciously evaluate the actions they might perform. In order to evaluate actions based on their probability of producing harm, then, individuals must have some awareness of those probabilities. But while the revealed-preference approach does show that an agent's behavior implicitly reflects a particular probability assignment, agents need not have any conscious awareness whatsoever of what probability they actually assign to any outcome: those probabilities describe agents' dispositions to choose, not their conscious thoughts. And because agents generally lack conscious access to the probabilities the theory ascribes to them, they will generally be unable to employ those probabilities in deliberating over whether to perform a particular action. Furthermore, while in theory an individual could herself employ the revealed-preferences approach to derive her own implicit probability estimates from her preferences, it is quite implausible that

---

<sup>53</sup> To be fair, some broadly descriptive approaches to decision theory often adopt a normative interpretation of decision theory's axioms—that is, the constraints on choice that an individual must meet if she is to be represented by a unique credence and utility function at all. Thus, though on this approach an individual may rationally assign any credence or any utility to any one particular proposition, rationality does impose certain constraints on the overall pattern of choices she may make. Nonetheless, criminal law is not concerned with this kind of guidance. The axioms of decision theory essentially impose a requirement of coherence on individuals: they rule out patterns of decision-making in which individuals implicitly rely on different probability estimates or different utilities in making different decisions. *E.g.*, De Finetti, *supra* note 49, at 7. Perhaps those axioms could themselves form reasonable guides to conduct, given that they generally describe direct constraints on permissible choices—*e.g.*, according to the sure-thing principle, an agent who would prefer one action to another both if a given proposition is true and if it is false should choose the one he prefers in both circumstances rather than the one he disprefers. But criminal prohibitions on recklessness do not merely require an agent's decisions to reflect consistent probability and utility judgments; rather, they require an agent to refrain from creating excessive risks to certain protected interests, which cannot be secured by compliance with the axioms of decision theory alone.

any real agent could actually employ that process in practice. Regardless of how exactly her probability assignments are understood to depend on her decision-making dispositions, actually identifying the probability assignment that would rationalize her choices would be complex and difficult: she must first (presumably introspectively) identify what those dispositions are, then analyze them mathematically to quantify the probability she associates with each outcome. Each step in this process is difficult in itself—individuals may not always know how they would decide in hypothetical situations, and complex mathematics may be required to derive quantified probabilities from decision-making dispositions. And even if she still could successfully complete them, performing this operation before every decision would be extremely impractical. Furthermore, this entire process presupposes that the agent’s hypothetical decisions are sufficiently coherent that her beliefs about risks can be represented by a single probability function. Empirically, however, agents systematically violate the constraints associated with standard expected-utility theory, and there is no agreement about whether and how the theory should be modified in response.<sup>54</sup> Thus, even if an agent could solve the practical challenges of actually applying this procedure in deliberation, there would be no guarantee in theory that her choices do reflect some single set of quantified probabilities that could be understood as her implicit estimate of the likelihoods of different outcomes.

Finally, even if an agent’s choices do reveal a unique, coherent set of quantitative probability assignments, those probabilities will likely be sufficiently inaccurate that individuals cannot rationally employ them to identify impermissibly risky conduct. The revealed preference approach does attribute implicit probability estimates to individuals even when they could not reason to those estimates based on their evidence. But if an agent’s evidence is insufficient to rationally justify explicit reasoning about quantitative probabilities that might be employed in evaluating actions, then any probabilities revealed by her choices cannot justifiably be employed in deliberation, either. Her preferences over hypothetical choices, after all, must depend on the same evidence that would justify her explicit reasoning about risk; if

---

<sup>54</sup> See, e.g., M. Allais, *Le Comportement de l’Homme Rationnel Devant le Risque: Critique des Postulats et Axiomes de l’École Américaine*, 21 *ECONOMETRICA* 503 (1953); Chandler, *supra* note 49, at 2–4; Daniel Ellsberg, *Risk, Ambiguity, and the Savage Axioms*, 75 *Q.J. ECON.* 643 (1961).

the latter is inaccurate, the implicit judgments guiding the former are no likelier to be accurate. Rather, whether those estimates are made explicitly through reasoning or revealed implicitly through judgments about hypothetical choices, they are likely to be little better than guesses. Because such guesses are unlikely to accurately identify the actions whose expected harms outweigh their benefits, deliberation employing those guesses will not effectively avoid such actions. Ultimately, then, even if agents could deliberate based on the probabilities revealed by their preferences, criminal prohibitions that require them to evaluate risky actions in deliberation based on those revealed probabilities will be an ineffective way of guiding individuals to avoid risk.

The law should not understand substantial and unjustifiable risk in terms of expected-utility reasoning, then, because quantitative probability estimates cannot plausibly play the role it assigns them in ordinary deliberation. A further reason also exists against understanding recklessness in terms of expected utility. In order to evaluate an action based on its expected utility, individuals must not only quantify the probability that it will cause various outcomes—which, I have argued, they generally cannot do accurately—but must also actually employ those probabilities in practice to compute expected utilities in deciding which actions to perform. And ordinary individuals making ordinary decisions can no more plausibly employ quantified probabilities in this way than they can identify those probabilities. First, were criminal law to guide conduct based on its expected utility, it would require individuals to value outcomes not according to their subjective valuations but rather according to objective values imposed by law.<sup>55</sup> But individuals are unlikely to accurately identify those objective utilities. Consequently, they are likely to incorrectly assess the expected utilities of actions. Second, even if individuals could accurately determine actions' expected utilities, it is simply unrealistic to expect them to do so routinely when making everyday decisions. Mental arithmetic takes time and effort, and humans simply lack the cognitive capacities to rely on it extensively in decision-making. Even if individuals could accurately quantify probabilities, then, they could not correctly identify which actions to avoid by computing expected

---

<sup>55</sup> *E.g.*, ALEXANDER, FERZAN, & MORSE, *supra* note 8, at 59; STARK, *supra* note 6, at 24.

utilities, both because they cannot realistically perform those computations and because those computations would not be accurate, if performed.

Prohibitions on recklessness aim to guide individuals to avoid actions that unjustifiably threaten others' interests, not actions that merely fail to advance an agent's own interests. Thus, if agents must evaluate the permissibility of creating particular risks based on the expected utility of doing so, they must employ utilities that capture the objective value of the harms and benefits an action causes to both the agent and others, not the subjective value the agent himself assigns to those outcomes. Furthermore, whether an action's expected utility is positive or negative depends as closely on the precise quantitative utilities assigned to those outcomes as on their precise quantitative probabilities. It is not enough to know that causing injury or death is bad; rather, whether an action's expected utility is positive or negative will depend on the relative value of causing injury and, say, making it to an important business meeting on time. Depending on how probable each outcome is, different utility values will yield different verdicts on the justifiability of the risk. Individuals' assessments of different actions' expected utilities will depend on whether they can accurately identify the objectively correct utilities of the outcomes those actions may cause. If they instead assign incorrect values to various possible outcomes of acting, their judgments as to whether the expected harms of an action outweigh its expected benefits will often be mistaken, and requiring them to act based on those judgments will not be an effective method of guiding them to avoid actions that in fact create an unjustified risk of causing harm.

But it is no more plausible that individuals can accurately identify the quantitative value of outcomes than that they can identify the quantitative probability that those outcomes will occur. Certainly, as a matter of psychological fact individuals do not ordinarily rely on quantitative utilities in reasoning (at least except in cases where outcomes themselves may be quantified): though we obviously have some sense of the values of different outcomes, we do not quantify them in order to reach a decision.<sup>56</sup> This failure to quantify the values of outcomes in actual

---

<sup>56</sup> *E.g.*, KNIGHT, *supra* note 15, at 68–69 (“It will be seen that the [utility] values of the curve have only the vaguest quantitative character. The boy not only does not ask how much sacrifice is how many berries worth, but



deliberation strongly suggests that individuals lack the ability to do so accurately. And, in fact, it is entirely unclear how individuals could identify the objective values of different outcomes. The law itself does not quantify them. While criminal sentences might quantify the disvalue of certain harms relative to one another, the expected utility of an action depends not on comparisons between harms but rather on the relative values of the harms and benefits that an action may cause. And while individuals do absorb the values of the societies in which they live, it is implausible that they could thereby learn the quantitative relationships between the values of different outcomes. Since in general the only things we can learn from other people are things that they actually know, only if ordinary members of our society themselves quantitatively valued outcomes could we learn those values from them. As I have argued, though, individuals do not ordinarily quantify utilities in making decisions, and thus we cannot learn how to value outcomes quantitatively from others' valuations. And it is implausible that such valuations could even be developed. Perhaps it might be possible to quantify the magnitude of certain highly salient harms—death, injury, and so forth—but because an action's expected utility depends on both its harms and its benefits, to compute expected utilities individuals would require a comprehensive valuation of every outcome that might result from acting. Given the detail and scope required of such a scheme, it is hardly surprising that none exist.

Finally, even setting aside that inaccuracies in individuals' assessments of expected utility would prevent them from effectively identifying excessively risky conduct, the effort required to actually compute expected utility in daily life would make unreasonable demands on individuals' cognitive capacities. Unlike many other kinds of law, which employ specialized rules to govern specialized fields of human activity, criminal law governs the ordinary decisions that individuals make throughout their daily lives. Individuals must be able to apply the rules of criminal law when making routine decisions. But computing the expected value of an action is an intensive process in terms of both time and cognitive resources: individuals must

---

merely, are these berries worth the sacrifice; he does not even ask, 'by how much' are these berries worth 'the' sacrifice. There is no true psychic quantity involved; only the commodity is measured or measurable.'").

first assign appropriate utility and probability figures to the different outcomes, then use those figures to calculate expected utilities, then use those expected utilities to decide how to act. Individuals could not feasibly employ this process in making each of the countless decisions they face over the course of an ordinary day: it is simply too taxing to be employed at such large scale.<sup>57</sup> Thus, individuals could not actually comply in deliberation with rules that require them to evaluate actions based on their expected utility. The rules for deliberation promulgated by definitions of recklessness in terms of quantified probabilities would not merely lead individuals to make poor decisions, were they employed; such rules cannot actually be employed at all in more than a small fraction of instances. And a rule that individuals cannot actually employ in deliberation does not merely guide conduct badly; it does not guide conduct at all.

To be sure, many decision theorists have recognized that precisely quantifying the likelihood of various outcomes is both “psychologically unrealistic” and “the wrong response to the sorts of evidence we typically receive.”<sup>58</sup> But they do not suggest abandoning quantitative models of decision-making altogether. Instead, they propose a different kind of quantitative representation of risk, on which an agent assigns sets of values to outcomes rather than single values. Each set, on this approach, contains any probability for that outcome that is compatible with the agent’s evidence.<sup>59</sup> Perhaps, then, definitions of recklessness could take a defendant’s awareness of risk to depend not on the single quantitative probability he assigns to an outcome but rather on the set of probabilities he assigns to it. Prohibitions on recklessness, thus defined, would still guide conduct by requiring individuals to evaluate actions by assessing the degree of risk they create; they would simply assess that degree of risk in a different way. Imprecise probabilities may provide a powerful theoretical account of reasoning under uncertainty. But they do not provide a

---

<sup>57</sup> E.g., SAVAGE, *supra* note 38, at 16 (“It is even utterly beyond our power to plan a picnic or to play a game of chess in accordance with the principle, even when the world of states and the set of available acts to be envisaged are artificially reduced to the narrowest reasonable limits.”). For further discussion of the costs of explicitly deliberating according to the formal models often employed to describe decision-making, see John Conlisk, *Why Bounded Rationality?*, 34 J. ECON. LIT. 669, 686–92 (1996).

<sup>58</sup> Joyce, *supra* note 39, at 283.

<sup>59</sup> See generally Seamus Bradley, *Imprecise Probabilities*, STAN. ENCYCLOPEDIA PHIL. (Edward N. Zalta ed., Spring 2019 ed.), <https://plato.stanford.edu/archives/spr2019/entries/imprecise-probabilities/>; Joyce, *supra* note 39; Alan Hájek & Michael Smithson, *Rationality and Indeterminate Probabilities*, 187 SYNTHESIS 33 (2012); Isaac Levi, *On Indeterminate Probabilities*, 71 J. PHIL. 391 (1974).

definition of recklessness that effectively guides individuals' behavior: if anything, individuals would avoid dangerous conduct even less successfully by employing imprecise probabilities to reason about risk than by employing precise ones.

First, though definitions of recklessness employing imprecise probabilities would require individuals to evaluate actions based on their imprecise likelihood of causing harm, it is not even clear what influence an imprecise measure of risk rationally should have on an individual's choices. A reasonably natural and intuitive decision rule exists for precise probabilities: individuals should decide what to do by discounting the value of the various outcomes by the likelihood that acting will produce them. But if probabilities are imprecise there exists no single value by which agents may discount those values. Instead, some alternative decision rule must evaluate actions using an imprecise credence interval rather than a precise credence value as an input. And unlike in the case of precise probabilities there is no natural, intuitive rule for making decisions with imprecise probabilities: though many candidates have been proposed and research continues, no consensus has emerged thus far. Seamus Bradley's conclusion from a recent survey is representative: "We have explored a number of different kinds of choice rule. None is entirely satisfactory. . . . All the more discriminating rules we have looked at seem to violate one or more intuitively compelling properties of rational choice."<sup>60</sup> It is possible that

---

<sup>60</sup> Seamus Bradley, *How to Choose Among Choice Functions*, in PROCEEDINGS OF THE 9TH INTERNATIONAL SYMPOSIUM ON IMPRECISE PROBABILITY: THEORIES AND APPLICATIONS 57, 64 (Thomas Augustin et al. eds., 2015). Bradley's reference to the "more discriminating" decision rules indicates that he accepts one very weak constraint on decisions made with imprecise credences—dominance, the principle that one act should be preferred to another just in case the former has greater expected value according to every single probability function compatible with one's evidence. *See id.* at 60. Others accept similarly weak principles: Paul Weirich, for example, defends a "permissive principle of choice" according to which an action may be chosen so long as it maximizes expected utility according to at least one element the set of probabilities compatible with the agent's evidence. PAUL WEIRICH, RATIONAL CHOICE USING IMPRECISE PROBABILITIES AND UTILITIES 54–58 (2021). But dominance alone provides very minimal constraints on choice: ordinarily the desirability of acting will differ considerably between the lowest and highest admissible estimates of a particular outcome. *See Bradley, supra*, at 62. Thus, it is implausible that the only requirement imposed by prohibitions on reckless conduct is that an agent not act when the expected value of acting is negative on any probability estimate that is compatible with the agent's evidence. Indeed, in some ways the decisions dominance would require of individuals seem to be exactly opposite of what the law should require. Because in general more evidence justifies more precise estimates of probability, the interval between the highest and lowest probability estimates compatible with an agent's evidence will be smaller the more evidence an agent possesses. Thus, as an agent gains more evidence dominance becomes likelier to prohibit certain actions. But it is not at all clear that this attitude towards risk is appropriate. If someone has very little evidence about the results of acting, dominance is unlikely to forbid any options, since so few probability functions will be eliminated by that evidence. But surely the absence of evidence is itself a good reason to refrain from acting: someone who does not know the likely results of action would seem to be reckless if he nonetheless acts.

some optimal decision rule for imprecise probabilities will be identified; it is also possible that there exists no uniquely rational method of evaluating actions under uncertainty unless risks can be precisely quantified.<sup>61</sup> But regardless of whether this problem will ultimately be solved, surely criminal law cannot today require ordinary individuals to use imprecise probabilities to evaluate actions when nobody can agree today how those probabilities rationally should affect an agent's choices.

Even if a decision rule for imprecise probabilities were to be found, furthermore, individuals would face the same problems in applying it that they do in evaluating actions using precise probabilities. In order to apply a decision rule that takes imprecise probabilities as inputs, an individual must know which probability interval, exactly, represents the likelihood of a particular outcome. For example, "take the midpoint," a very crude decision rule that instructs individuals to compute expected utilities using the midpoint of each probability interval, requires an individual to know the upper and lower bound of the interval that represents each outcome in order to find the midpoint between them.<sup>62</sup> As with precise probabilities, then, to evaluate actions using imprecise probabilities individuals must be able to explicitly quantify their beliefs about the likelihood of each outcome; the only difference is that they must precisely quantify an interval rather than a point value. But explicitly quantifying the interval of probabilities compatible with one's evidence is no easier than explicitly quantifying a single value that evidence uniquely identifies. In general, probability judgments represent the proportion of possible states of affairs in which some outcome obtains.<sup>63</sup> Imprecise probabilities might be easier to quantify than precise ones, then, if it were easier to identify upper and lower bounds on that proportion than to identify it exactly. But our evidence is typically inadequate to precisely quantify probabilities because it does not enumerate some set of possibilities that we could examine to determine how frequently a particular outcome occurs. And without some method of analyzing such possibilities, we can no more identify the range in which the frequency of that outcome must fall than we can identify it exactly. Because individuals will ordinarily lack explicit beliefs about the

---

<sup>61</sup> For some skepticism, see KEYNES, *supra* note 15, at 361; and Bradley, *supra* note 60, at 64–65.

<sup>62</sup> See Joyce, *supra* note 39, at 311–12.

<sup>63</sup> See *supra* notes 40–42 and accompanying text.

quantitative intervals that represent the probabilities of particular outcomes, then, they cannot apply a decision rule that takes those intervals as inputs.

Finally, employing imprecise probabilities would exacerbate rather than alleviating the demands that expected-value calculations place on human cognitive capacities. Agents who represent probabilities with single values must apply a relatively straightforward decision rule that sums the utilities of the possible outcomes of an action, each discounted by its probability of occurring. But a rule that evaluates actions based on intervals or sets rather than single values must be more complex to accommodate the greater complexity of using intervals or sets as an input rather than single values.<sup>64</sup> A more complex decision rule would be even more difficult to apply. Thus, if individuals making routine decisions cannot plausibly compute expected utilities using precise probabilities, it would be even less plausible for them to compute expected utilities using whatever decision rule might be developed to reason with imprecise probabilities.

To be sure, I do not mean to claim that expected-utility reasoning is never possible or useful. The routine decisions that arise in everyday life do not present individuals with sufficient evidence to rationally justify quantifying the likelihood that a particular outcome will occur, nor could individuals realistically compute expected utilities when making such decisions. But the fact that expected-utility reasoning cannot always be employed in decision-making obviously does not entail that it cannot be employed in decisions that do not feature the obstacles that prevent its use in everyday contexts. Because an administrative agency can devote considerable manpower and time to rulemaking, for example, it may well be able to develop or identify some empirical basis for quantifying the probable outcomes of various possible legislative rules. Expected-utility reasoning may well be feasible in those contexts. Indeed, explicitly calculating expected utilities may be the optimal form of decision-making under uncertainty. But even if it is in some sense optimal, the conditions that make its use possible are met only rarely. In particular, they are almost never met in the ordinary decisions that ordinary individuals must constantly make when going about their lives—which are precisely the kinds of decisions that

---

<sup>64</sup> See generally Bradley, *supra* note 60, for an overview of some rules that could be employed.

criminal law must regulate. Thus, whatever the importance of expected-utility reasoning in some contexts, criminal law cannot effectively guide conduct through rules that identify prohibited actions in terms of their expected utility, since individuals cannot accurately apply those rules in more than a small fraction of the decisions they actually face.

Defining recklessness in terms of explicit attitudes towards risk requires individuals to evaluate actions based on whether their likelihood of causing harm exceeds some threshold of substantiality and unjustifiability. As Alexander, Ferzan, and Morse explain, on their view the prohibition on recklessness requires individuals to “choose only those acts for which the risks to others’ interests – as you estimate those risks – are sufficiently low to be outweighed by the interests, to yourself and others, that you are attempting to advance (discounted by the probability of advancing those interests).”<sup>65</sup> So long as the key notions in this rule, such as weighting and discounting, are understood in their most natural, quantitative sense, then this injunction is—as Savage says—preposterous.<sup>66</sup> Because we cannot accurately identify the appropriate weights and discounts, and because we cannot plausibly compute the expected utility of every action we consider, we cannot effectively avoid dangerous actions by following this rule. In some sense the rule might provide an accurate specification of wrongful acts: individuals should in fact avoid the conduct that it declares to be impermissible. But since criminal prohibitions aim to guide conduct, they cannot employ just any specification of wrongful acts; rather, they must identify a decision procedure that individuals can actually employ in deliberation.<sup>67</sup> And because a rule evaluating actions based on quantified probabilities does not effectively guide individuals to avoid causing harm, criminal prohibitions should not define offenses so as to require individuals to employ rules of that sort in deliberation.

---

<sup>65</sup> ALEXANDER, FERZAN, & MORSE, *supra* note 8, at 263.

<sup>66</sup> SAVAGE, *supra* note 38, at 16.

<sup>67</sup> For an explanation of the distinction between a decision procedure and a criterion of right action, see, for example, Peter Railton, *Alienation, Consequentialism, and the Demands of Morality*, 13 *PHILOSOPHY & PUBLIC AFFAIRS* 134, 152–53 (1984).

### *B. Qualitative Reasoning About Risk*

Individuals cannot plausibly guide their conduct by explicitly quantifying the degree of risk their actions would create. But perhaps individuals could nonetheless reason explicitly about risk in ordinary deliberation by employing some sort of non-quantitative representation of risk—that is, perhaps individuals could ordinarily assess the risk individual actions create in some qualitative sense even if they cannot measure that risk numerically. This understanding of how individuals reason about risk might then be employed to develop an account of when defendants are aware of risk. Defendants would still possess the sort of awareness of risk required for recklessness only if they possess some explicit mental attitude directed towards the degree of risk that an action would create, but that attitude would represent the degree of risk created qualitatively rather than quantitatively. But even on a qualitative interpretation of awareness of risk, not all risk-creation would be prohibited; thus, criminal prohibitions could guide conduct only by somehow identifying in qualitative terms which risks are prohibited. Recklessness is the conscious disregard only of a substantial and unjustifiable risk; if individuals are to guide their deliberation based on their qualitative assessments of risk, then, they must be able to identify whether a particular risk, given its qualitative degree, is substantial and unjustifiable. And, I will argue, the law cannot promulgate a rule for individuals to employ that articulates some qualitative threshold that individuals could employ in evaluating actions. Consequently, criminal prohibitions cannot effectively guide individuals' conduct by defining recklessness in terms of a mental state that qualitatively represents the degree of risk an action creates.

Qualitative reasoning about risk requires qualitative assessments of the degree of risk that particular actions create. How, then, might individuals assess risk qualitatively? One suggestion might look to the sorts of qualitative judgments that individuals do in fact make in reasoning about risk. In deliberation, we do note that certain risks are high or low, or that certain harms and benefits are serious or slight, even though we would not be able to translate those judgments into some numerical scale on which we could measure the quantitative difference between any probability

ranging from impossibility to certainty.<sup>68</sup> Perhaps these sorts of qualitative assessments should determine whether a particular defendant was reckless: to be aware of a substantial and unjustifiable risk, a defendant must assess the risks an action creates to exceed some threshold defined in terms of qualitative judgments, such as judgments about whether the risk an action creates is high or low.<sup>69</sup> Defining recklessness in this way, then, would require defendants to evaluate actions based on whether they create a risk greater than that threshold.

Quantitative risk assessments cannot effectively guide conduct, I have argued, because individuals' evidence does not ordinarily justify assigning a single probability to various outcomes that might result from action. But while individuals do not ordinarily quantify risk in everyday decisions, qualitative risk assessments are routine. In this respect, they do avoid an obstacle facing quantitative ones; nonetheless, in a different respect they fall prey to an obstacle that quantitative assessments avoid. While quantified probabilities cannot effectively guide conduct in practice, it is reasonably clear in theory how those estimates should influence choice, could they be made. By contrast, while in practice qualitative risk assessments may be available in ordinary deliberation, no account exists in theory of how they rationally ought to influence choice. Expected-utility reasoning provides a clear rule specifying whether an action may be performed given its quantitative probability of causing various results. But no rule determines whether an action may be performed given the qualitative degree of risk it will create. May an agent perform an action that has a low risk of causing a severe harm but a high risk of causing a mild benefit, for example? Obviously, no answer can be given at this level of generality; the right choice depends on how high and low the risks are, and how mild and severe the harm and benefit are. Whether harms outweigh benefits depends on how much worse the harms are than the benefits; how much to discount the value of an uncertain outcome depends on how likely it is not to occur. And to reach an overall verdict on an action one must be able to aggregate all these separate comparisons into a single judgment

---

<sup>68</sup> For Keynes, numerical measurements of likelihood have two distinguishing features: they employ a scale that ranges from certainty to impossibility, and the values on the scale differ in magnitudes that may be manipulated using the rules of arithmetic. *See* KEYNES, *supra* note 15, at 180–82.

<sup>69</sup> Alexander, Ferzan, and Morse suggest this possibility. ALEXANDER, FERZAN, & MORSE, *supra* note 8, at 282.



as to whether, given how much worse the harms are than the benefits, and how much likelier the benefits are to occur, the action may be performed. Because the laws of arithmetic apply to quantitative probabilities, they allow quantitative risk assessments to be aggregated into an overall judgment. But because qualitative assessments do not obey arithmetic rules, no such aggregation procedure can be applied to them.

Ultimately, then, qualitative risk assessments cannot guide conduct effectively for the same reason that quantitative assessments cannot. Expected-utility reasoning does state a plausible criterion for whether an action may be performed given its likelihood of causing various possible outcomes: the actions we should avoid in fact are those whose expected value is sufficiently negative. For individuals themselves to apply this criterion in deliberation, though, they must be able to identify the expected values of actions, and as I have argued the available evidence does not ordinarily permit them to do so. The very same problem confronts qualitative risk estimates. To evaluate actions through explicit reasoning about risk, agents must employ their evidence to evaluate the expected value of their actions precisely enough to determine whether those actions may be performed. But if the evidence is inadequate to reliably assess actions' expected value, it will be inadequate regardless of how individuals reason about risk; different kinds of reasoning will simply locate the unreliable inference in different places. Quantitative reasoning about risk allows agents to reliably infer whether they may act given how risky an action is, but they cannot reliably infer how risky an action is given their evidence about it; qualitative reasoning allows agents to reliably infer how risky an action is given their evidence about it, but they cannot reliably infer whether an action may be performed given how risky it is.<sup>70</sup> Since both steps are necessary to reliably identify which actions are too dangerous to be performed, individuals cannot employ approach to reasoning about risk in order to effectively guide their conduct.

Furthermore, it is not even clear how recklessness could be defined using a qualitative understanding of the magnitude of risk. As I have argued, criminal

---

<sup>70</sup> To be fair, it might be possible to define qualitative risk assessments so that they entail certain evaluative conclusions about actions: for example, perhaps the judgment that a risk is high should entail that an action may not be performed. Certainly, such a definition would be possible, but it would simply replace uncertainty about how to act given certain qualitative risk assessments with uncertainty about which risks should be assessed as high.

prohibitions guide individuals to employ rules in deliberation by defining offenses to encompass actions chosen in violation of those rules. Defining recklessness so that whether a defendant is reckless depends on the degree of risk he believed or intended an action to create, then, promulgates a rule that requires individuals to evaluate actions based on the degree of risk they believe or intend those actions to create. For example, definitions of recklessness employing quantified probabilities require individuals not to act when an action would have negative expected utility given its quantitative probability of causing harm. But while this rule identifies actions individuals may not perform based on their quantitative probability of causing harm, no analogous qualitative rule exists. As I have argued, risk-creation is not forbidden simply because the risk is high, or because the harm risked is severe; qualitative risk assessments do not on their own determine whether individuals may act. And if no qualitative judgment that individuals could make about an action's degree of risk would require them not to perform it, what qualitative judgment about its degree of risk would make them culpable for performing it? If a defendant must believe or intend his action to create a degree of risk that is in fact substantial and unjustifiable for him to be reckless, then the law must somehow identify what degrees of risk are in fact substantial and unjustifiable, for only defendants with mental states directed towards risks of that degree would be guilty of recklessness. But precisely because no rule specifies qualitatively which risks may not be created, substantial and unjustifiable risk cannot be defined in qualitative terms. Thus, attempts to define recklessness in terms of a qualitative awareness of risk cannot characterize what kind of qualitative awareness of risk a defendant must possess for him to be guilty of recklessness.

Individuals do not reason about risk by applying some rule that evaluates actions based on explicit qualitative assessments of the degree of risk they would create. Consequently, the awareness of substantial and unjustifiable risk required for recklessness cannot be defined as the qualitative attitude towards risk from which an individual must conclude, in deliberating, that acting would be forbidden. Nonetheless, I do not mean to deny the obvious point that individuals clearly have some capacity to evaluate the riskiness of particular actions. We must have some such

capacity, after all, since we do sometimes make decisions when the results of acting are not known. And it is easy to think of examples of actions for which the appropriate risk assessment is obvious: everyone knows that shooting a gun in the air in Times Square would be too risky, say, or car racing during rush hour down a crowded city street. And since somehow individuals can qualitatively assess those actions accurately enough to know they are too risky, perhaps recklessness could be defined in terms of whatever assessments of risk are employed in that process. Indeed, I will generally defend an account of recklessness along these lines. But the method individuals actually do employ to evaluate actions requires a deeper reconceptualization of the kind of deliberative process that prohibitions on recklessness aim to regulate: because individuals do not evaluate actions by attempting to directly assess the magnitudes of individual risks, prohibitions on recklessness cannot guide conduct by identifying which risks are substantial and unjustifiable in terms of their magnitude.

Certainly, it is not difficult to imagine situations, from the everyday to the momentous, in which individuals must decide what to do despite uncertainty about the outcomes of acting. We do not typically make these decisions, I have argued, by first assessing the likelihoods that acting will lead to various outcomes, then deducing the correct decision from those assessments. How, then, do we decide? Broadly speaking, we act by judging risk holistically: in Keynes' words, we make decisions based on "an intuitive judgment directed to the situation as a whole, and not in virtue of an arithmetical deduction derived from a series of separate judgments directed to the individual alternatives each treated in isolation."<sup>71</sup> Similarly, Knight emphasizes "that the exact science of inference has little place in forming the opinions upon which decisions of conduct are based . . . . We act upon estimates rather than inferences, upon 'judgment' or 'intuition,' not reasoning, for the most part."<sup>72</sup> And, crucially, that intuition is based on a holistic assessment of our choices and their context: "[I]n the main, it seems that we 'infer' largely from our experience of the past as a whole, somewhat in the same way that we deal with intrinsically simple

---

<sup>71</sup> KEYNES, *supra* note 15, at 357.

<sup>72</sup> KNIGHT, *supra* note 15, at 223.

(unanalyzable) problems like estimating distances, weights, or other physical magnitudes . . . .”<sup>73</sup> I do not mean to speculate about the psychological basis of such judgments; as Knight puts it, “The ultimate logic, or psychology, of these deliberations is obscure, a part of the scientifically unfathomable mystery of life and mind.”<sup>74</sup> Perhaps they involve a skilled capacity to evaluate risk that cannot be translated into conceptual terms; perhaps they are merely essentially random guesses that we employ when no better methods are available.<sup>75</sup> But by and large, these intuitive evaluations form the basis of our actual decisions about how to act when acting would be risky.

If individuals actually employ this sort of procedure in deciding whether or not to create certain risks, prohibitions on criminal recklessness must promulgate a rule governing this sort of decision-making in order to guide conduct. Indeed, in this chapter I will defend an account of criminal recklessness on which it defines the mental state defendants must possess to be reckless so as to promulgate a rule that governs this process. But for criminal recklessness to guide this sort of holistic judgment about risk would require a very different conceptualization of the mental state that is required for recklessness than the one employed on interpretations of recklessness as intentional risk-creation. In particular, those interpretations of recklessness understand the mental state that determines whether defendants are reckless to be one whose object is the degree of risk created by a particular action: if recklessness requires a defendant to be consciously aware of a substantial and unjustifiable risk, then since the magnitude of a risk determines whether it is substantial and unjustifiable, a defendant’s awareness must extend to the magnitude of that risk in order to distinguish between defendants who are and are not reckless. This definition of recklessness, I have argued, instructs individuals how to act given their expectations concerning the likelihood that acting will cause various outcomes:

---

<sup>73</sup> *Id.* at 211.

<sup>74</sup> *Id.* at 227.

<sup>75</sup> For contrasting perspectives on the question, compare *id.* (“We must simply fall back upon a ‘capacity’ in the intelligent animal to form more or less correct judgments about things, an intuitive sense of values. We are so built that what seems to us reasonable is likely to be confirmed by experience, or we could not live in the world at all.”); with KEYNES, *supra* note 15, at 366 (“Thus the philosopher must draw what comfort he can from the conclusion with which his theory furnishes him, that millionaires are often fortunate fools who have thriven on unfortunate ones.”).

in particular, individuals may not perform any action that creates risks whose magnitude exceeds the threshold established by the definition of substantial and unjustifiable that the law employs.

Central to this account of both when individuals are reckless and how risk should influence their decision-making, though, is the presupposition that individuals do actually form mental states directed towards the magnitude of various risks when deciding how to act. By contrast, I have argued that the method individuals actually employ to make decisions need not feature any such mental states. Individuals normally decide what to do by intuitively assessing the reasonableness of the risks that would be created by performing a particular action for a particular purpose in a particular context. Crucially, this method evaluates actions and their circumstances context holistically; we intuitively reach a judgment as to the reasonableness of an action by directly considering all its relevant features. Of course, in holistically evaluating an action we will no doubt consider the sorts of harm that it may cause, as well as the reasons why a particular risk of harm is more or less serious. But what makes such evaluations holistic is that they are based directly on the underlying facts about an action. We do not first form some intermediate judgment about the magnitude of each risk our action might create, then employ those intermediate judgments to form a final judgment about whether to perform the action; instead, deliberation moves directly from facts about an action and its circumstances to a conclusion about what to do. We know that shooting a gun in Times Square is too risky, say, by directly and holistically evaluating the action itself, not by first computing the degree of risk it creates, then evaluating it based on that computation. And if in reasoning about risk our conclusions about whether to perform an action follow directly from our beliefs about the nature and circumstances of that action, the rules criminal prohibitions promulgate must guide that sort of inferential step: they must instruct individuals how to identify which actions they may not perform based on their beliefs about the nature of an action and its circumstances.

Interpretations of recklessness as intentional risk-creation, then, ultimately rest on an untenable account of how individuals reason about risk. Such interpretations take recklessness to require individuals to possess an intention or belief that the action

they perform creates a risk of a particular degree. By defining acting with that mental state to be an offense, then, this approach provides individuals with a rule specifying how their assessments of the degree of risk their actions will create should influence their decisions about which actions to perform: in particular, individuals must choose based on whether the degree of risk they take their actions to create exceeds the degree of risk defined to be substantial and unjustifiable. Such rules instruct individuals whether to act given an explicit judgment about how risky acting would be. Rules of this sort would guide conduct effectively if individuals in fact did reason about risk by explicitly judging how risky different actions would be. But we do not reason about risk in that way, and no matter how degrees of risk are conceptualized we should not reason about risk in that way. Rules that instruct individuals how to reason about risk in this manner simply govern conduct ineffectively: the law cannot effectively reduce actions that create unjustified risks by requiring individuals to employ a rule in deliberation that they cannot employ to accurately identify which actions will, in fact, create unjustified risks. Instead, in order to effectively guide individuals to avoid actions that create unjustified risks, the law must promulgate rules that effectively govern the process individuals actually do employ to reason about risk.

## **II. Acting Under Uncertainty**

If prohibitions on recklessness are to effectively guide individuals to avoid excessively dangerous actions, they must identify a rule for individuals to follow that can be applied to govern the methods individuals actually employ to reason about risk. Though it might be natural to think that individuals reason about risk by explicitly considering the degree of risk different actions create, I have argued that the evidence available is insufficiently detailed to justify any specific, explicit conclusions about how risky a particular action would be. Those explicit representations of the degree of an action's risk, then, cannot reliably identify which actions are too dangerous to perform. Instead, I have suggested, individuals reason about risk through a holistic evaluation of all the information they possess concerning a particular action they might perform. This suggestion, then, adopts a different model of how individuals proceed from their premises to their conclusion when

deliberating over what to do. On the view that I have criticized, that deliberation is essentially a two-step process. Individuals begin with various beliefs and intentions about a possible action and about the circumstances in which they would perform it. On the basis of that information, they judge the likelihood that the action would cause various results. Then, using those assessments of likelihood, they determine whether or not to perform the action. The rules promulgated by prohibitions on recklessness, on this approach, guide the second step in deliberation: they instruct individuals what to do based on how they assess the likelihood that an action will cause harm. By contrast, the holistic approach to judging risk that Keynes and Knight suggest individuals employ in actual deliberation is a one-step process: individuals begin with their beliefs and intentions about their action, then decide directly on the basis of that information. Rules that guide this process must answer a different question—namely, how to act given one’s total information about an action and its circumstances.

In this part, I will defend an account of how criminal prohibitions on recklessness should answer this question. As I have noted, the Model Penal Code and other legal definitions of recklessness often require defendants to be aware of risk to be convicted.<sup>76</sup> Many scholars argue that recklessness must involve explicit mental states directed towards the degree of risk an action creates by interpreting “awareness of risk” to refer to such a mental state.<sup>77</sup> And since recklessness requires defendants to be aware of risk, reckless conduct therefore must involve some explicit attitude towards risk. First, then, I will argue for a reconceptualization of the notion of awareness of risk that recklessness employs. Recklessness, in my view, involves what probability theorists have called uncertainty rather than risk: to be aware of risk agents need not explicitly recognize some likelihood of harm but rather must simply lack knowledge of the results of acting.<sup>78</sup> Recklessness, then, guides conduct by providing a rule for individuals to follow when they do not know what results their actions will cause. What rule, then, should that be? The central flaw in interpretations of recklessness as intentional risk-creation, I argued, was that individuals lack adequate evidence to reliably determine, on their own, which actions are too

---

<sup>76</sup> See *supra* note 23.

<sup>77</sup> E.g., STARK, *supra* note 6, at 92–94; Ferzan, *supra* note 8, at 636–38; Husak, *supra* note 7, at 208.

<sup>78</sup> I am grateful to Gideon Yaffe for suggesting that the risk/uncertainty distinction might be fruitful for me.

dangerous to be performed. How should individuals act if they cannot reliably identify dangerous conduct on their own? They should defer to the judgment of others. Prohibitions on recklessness, in my view, promulgate this rule: they require individuals to avoid actions that law-abiding members of their society would not perform.

### *A. Risk and Uncertainty*

To be reckless, a defendant must be aware of risk. One obvious way to interpret the notion of awareness of risk is to understand it in the same way that we understand awareness of anything else. To be aware of the cat lying next to one's feet, one must have some mental state directed towards it—one must see it, or perceive it, or simply believe it to be there (if, say, one has been told it is there). Broadly speaking, to interpret recklessness to require some mental state directed towards risk is to adopt this approach: individuals who act believing that a particular risk exists are aware of that risk in the same sense that someone who sees the cat is aware that it exists. I have criticized these interpretations on normative grounds, since the rule of conduct they promulgate for individuals to employ in deliberation does not effectively guide them to avoid dangerous actions. But the law itself does require recklessness to involve the awareness of risk.<sup>79</sup> Thus, if the interpretation of awareness of risk in terms of beliefs directed towards risk should be rejected on normative grounds, I must show how an alternative understanding of recklessness can still vindicate the law's requirement that reckless defendants be aware of risk.

The familiar culpability notions employed by the Model Penal Code form a hierarchical structure: purpose, knowledge, recklessness, and negligence each define a form of culpability less serious than the one above it in the hierarchy.<sup>80</sup> Knowledge, the form of mens rea immediately superior to recklessness, requires a defendant to be “practically certain that his conduct will cause . . . a result” for him to be culpable for causing it.<sup>81</sup> The mental state required for recklessness, then, must in some way fall short of practical certainty. To be sure, a belief or other mental state directed towards

---

<sup>79</sup> MODEL PENAL CODE § 2.02 cmt. 3 (AM. LAW INST. 1962).

<sup>80</sup> See MODEL PENAL CODE § 2.02(5) (noting that each form of culpability is automatically satisfied if a defendant satisfies the requirements of a form of culpability located higher in the hierarchy).

<sup>81</sup> MODEL PENAL CODE § 2.02(2)(b)(ii).



a risk of a particular degree does in some way fall short of practical certainty: someone who believes there to be a particular degree of risk that acting will cause some result lacks certainty as to whether acting will in fact cause that result. But in another sense, someone who believes that an action has some precise likelihood of causing harm does possess practical certainty, just of a different sort. The person who flips a coin, say, does not know whether it will come up heads or tails, but he does know what the chance of each outcome is. He does not lack knowledge; rather, he merely possesses a different, probabilistic, kind of knowledge, which concerns likelihoods rather than states of affairs.<sup>82</sup> And individuals may quite rationally take an attitude of practical certainty towards information about the likelihoods of outcomes, just as they may towards information about outcomes themselves, by relying on the truth of that information in deliberating about what to do. Someone facing a decision affected by the flip of a fair coin should take for granted the equal likelihood of heads and tails in evaluating possible actions, just as she would take many other facts for granted in deliberating about what to do. Such risks involve only partial uncertainty—uncertainty about the actual results of acting coupled with certainty about the probability of various possible results.

Consequently, many theorists of probability—notably, both Keynes and Knight—have argued that there are in fact two different ways in which individuals making decisions can lack knowledge of the results of possible actions.<sup>83</sup> First, in some decisions no possible outcome of acting is practically certain, but the likelihoods of different outcomes are measurable. In Knight’s terminology, these decisions present agents with risks.<sup>84</sup> In this technical sense, then, risks exist only when the evidence available to an agent in deliberation permits the probability of different outcomes to be numerically quantified. But as both Knight and Keynes argue, such evidence is only rarely accessible to individuals in deliberation.<sup>85</sup> In most contexts, then, individuals instead face what Knight calls “true uncertainty,” in which

---

<sup>82</sup> For an overview of this sort of knowledge, see generally SARAH MOSS, *PROBABILISTIC KNOWLEDGE* (2018).

<sup>83</sup> For more on Keynes and Knight’s role in introducing the distinction, see generally Phil Faulkner et al., *F. H. Knight’s Risk, Uncertainty, and Profit and J. M. Keynes’ Treatise on Probability after 100 years*, 45 *CAMBRIDGE JOURNAL OF ECONOMICS* 857 (2021), and the papers discussed therein.

<sup>84</sup> KNIGHT, *supra* note 15, at 233 (associating “risk” with “measurable uncertainties or probabilities”).

<sup>85</sup> See *supra* notes 43–48 and accompanying text.

they lack certainty about both what the result of acting will be and how likely various possible results are.<sup>86</sup> Agents face uncertainty when their evidence fails to rationally justify any expectation about what a particular action will cause. In such cases, as Keynes put it, “there is no scientific basis on which to form any calculable probability whatever.”<sup>87</sup> Instead, “[w]e simply do not know” what such an action will cause.<sup>88</sup>

Both true uncertainty and risk in this technical sense are plausible understandings of what it might mean for an agent to be aware of risk. But the mental states involved in each are quite different. A risk in the technical sense involves a particular measurable likelihood that acting will produce a result. Because that measurable probability must exist for the agent to face risk rather than uncertainty, the agent must be consciously aware of that probability to be consciously aware of the risk—as, for example, by believing, knowing, or intending a risk of that degree to exist. Agents face risk when their evidence rationally determines the probability of harm, and they are aware of that risk by recognizing, through the possession of some appropriate mental state, the probability of harm rationally determined by the evidence. But while risk involves a particular measurable probability, uncertainty involves the absence of any measurable probabilities. An agent is uncertain simply when, given his evidence, he lacks knowledge of either the actual or probable results of acting. And to lack knowledge is not to possess any particular mental state but rather not to possess one.<sup>89</sup> An uncertain defendant is aware of risk, then, simply by being uncertain—his awareness of risk lies in his lack of practical certainty concerning the actual or probable results of acting. This sort of awareness does not

---

<sup>86</sup> KNIGHT, *supra* note 15, at 232–33.

<sup>87</sup> J.M. Keynes, *The General Theory of Employment*, 51 Q.J. ECON. 209, 214 (1937).

<sup>88</sup> *Id.*

<sup>89</sup> Others have suggested that the awareness of risk required for recklessness might be satisfied not only by the existence of some positive mental state but rather also by its absence: Glanville Williams, for example, argued that the requirement of awareness of risk is satisfied by a defendant who lacks the belief that no risk exists. *See* Williams, *supra* note 8, at 84. Nonetheless, my suggestion differs from Williams’s in two respects. First, as should be obvious, mere uncertainty (whether construed probabilistically or as the absence of full belief, as on Williams’s approach) cannot on its own suffice for recklessness, since all decisions are made under uncertainty as to some aspects of an action’s results. Thus, uncertain defendants must satisfy some further condition to be reckless. Williams, however, identifies no such further condition. Second, Williams understands uncertainty as the lack of full belief, whereas I understand it as the absence of a precise estimate of the likelihood of harm. In my view, whether an agent fully believes that a risk exists in part reflects the agent’s evaluative judgments about which risks are reasonable to disregard. Thus, in my view an agent who fully believes that a risk exists may nonetheless be culpable, since he may have culpably disregarded the risk by forming the belief that it did not exist. *See infra* Section I.C.

require the existence of any particular mental state, since uncertainty cannot be reduced to a single mental state. But uncertainty remains a matter of individuals' conscious awareness, for whether an individual is uncertain depends on what an individual is consciously aware of; it differs from awareness of risk in the technical sense only in that the defendant is aware of uncertainty because a particular mental state is absent from his conscious awareness, not because a particular mental state is present in it.

Which of these two notions of risk should be employed in understanding the notion of awareness of risk, as it is employed in the law?<sup>90</sup> Most commentators appear to find an interpretation of awareness of risk in terms of beliefs about risk to be obvious and natural, devoting little argument to the claim that awareness requires belief or knowledge.<sup>91</sup> But though the step from awareness to belief has seemed obvious to commentators, in fact it receives little support from the law itself. The text of the Model Penal Code, as I have noted, only vaguely specifies the mental state required for recklessness. And its commentaries conspicuously fail to define awareness of risk in terms of a belief explicitly about risk. They gloss "awareness of risk" in two ways. First, they specify that awareness is required "of a probability less than substantial certainty."<sup>92</sup> A probability less than substantial certainly is clearly a form of uncertainty. But the commentaries are here ambiguous between measurable and unmeasurable forms of uncertainty: though they gloss a risk as a probability less than substantial certainty, they do not specify whether that probability must be numerically measurable, nor do they specify whether defendants must be able to specify quantitatively how much less than certain that risk is. The next clause in that sentence, furthermore, provides a description of when an agent is aware of risk that seems clearly to equate it with true uncertainty: awareness of risk exists when "the matter is contingent from the actor's point of view."<sup>93</sup> This definition is clearly broad

---

<sup>90</sup> The Model Penal Code's use of the word "risk," which obviously aligns terminologically with a particular side of the risk/uncertainty distinction, might seem at first glance to favor one interpretation of the notion of awareness of risk. But I am aware of no reason to think that the drafters of the code sought to track this particular technical term of art from probability theory rather than simply using "risk" in its ordinary sense, which Knight himself agrees encompasses the technical notions of both risk and true uncertainty. KNIGHT, *supra* note 15, at 233.

<sup>91</sup> *E.g.*, STARK, *supra* note 6, at 92–93; Ferzan, *supra* note 8, at 597–98; Husak, *supra* note 7, at 208.

<sup>92</sup> MODEL PENAL CODE § 2.02 cmt. 3 (AM. LAW INST. 1962).

<sup>93</sup> *Id.*

enough to encompass both measurable and unmeasurable uncertainty: the outcome of acting is contingent from an individual's point of view simply when, given what that individual is aware of, no outcome is necessary. Notably, this definition does not require the agent to believe that the matter is contingent; the matter must only be contingent. Thus, contrary to the assumptions of most commentators, the code's own commentaries interpret awareness of risk not to require any particular belief about the risk an action will create: so long as an agent's evidence does not make the outcome of acting practically certain, the agent is aware of a risk.

My interpretation of awareness of risk, then, frames prohibitions on reckless conduct as answering a very different kind of question about how individuals must act. Defendants are reckless only if they are aware of risk. If the awareness of risk involves an explicit belief about the degree of risk an action will create, then, prohibitions on recklessness promulgate a rule instructing individuals what to do when they have a particular belief about the degree of risk an action would create—in particular, it tells them not to act if the risk they believe that action would create is substantial and unjustifiable. By contrast, on my interpretation individuals are aware of risk when they simply do not know what will happen if they act—when, as the commentaries put it, the results of acting are contingent from their point of view. And prohibitions on recklessness, then, instruct individuals how to act in those circumstances—what to do when you don't know what's going to happen. Obviously, though, if individuals do not have any belief directed towards the magnitude of the risk an action would create, they cannot decide what to do based on how risky an action would be. How, then, should they decide instead? What rule should individuals follow when they don't know the outcome of acting?

### *B. The Law-Abiding Person*

As I have noted, expected-utility reasoning does in some sense correctly distinguish permissible from prohibited risk-creation: individuals should, in fact, avoid creating risks whose expected harms are not justified by their benefits. But though this criterion does identify which actions individuals should perform, it cannot be applied accurately in deliberation. Thus, though individuals should strive to avoid performing unjustifiably dangerous actions, they must instead deliberate using some

alternative procedure to identify whether any particular action is actually unjustifiably dangerous. The problem of identifying workable decision procedures extends beyond issues of risk, however. My arguments against interpreting recklessness in terms of attitudes about risk claimed that explicit expected-utility reasoning is not a feasible method for individuals to employ generally in making ordinary decisions. And, empirically, the results of decision-making often diverge from what expected-value maximization would require.<sup>94</sup> How, then, do individuals make decisions? Broadly speaking, individuals decide by following heuristics or rules of thumb, quick and easy methods of deciding what to do when the consequences of acting cannot be accurately anticipated and evaluated and thus cannot form a basis for choice.<sup>95</sup> If individuals generally deliberate and decide using such rules, then the law may guide them to avoid dangerous actions by identifying a heuristic individuals must employ when making everyday decisions. If both individuals' evidence and their cognitive capacities are inadequate for reliable evaluations of the expected results of acting, individuals may instead make decisions effectively by following some rule that can be easily applied in routine contexts and that will identify accurately enough the actions that must be avoided because of their risk.

Since this rule, as I have framed it, must address individuals who face uncertainty, rather than risk in its technical sense, one general type of rule cannot be employed—namely, rules that distinguish between different risks based on their likelihood of causing harm. Nonetheless, the degree of the risks an action creates is obviously an enormously important factor affecting whether it should be performed. When individuals act under uncertainty, then, a key factor normatively relevant to their choices cannot affect their decisions, precisely because—due to their uncertainty—they do not know the facts they ideally would employ in choosing. But if uncertainty undermines individuals' ability to rationally evaluate actions, then uncertainty itself would appear to be rationally relevant to choice. Individuals are less

---

<sup>94</sup> See Conlisk, *supra* note 57, at 669–75 for a survey of the literature.

<sup>95</sup> The importance of such rules has long been recognized: Mill, for example, explained that such rules “point out the manner in which it will be least perilous to act, where time or means do not exist for analysing the actual circumstances of the case, or where we cannot trust our judgment in estimating them.” 8 JOHN STUART MILL, COLLECTED WORKS OF JOHN STUART MILL 946 (J.M. Robson ed., 1974). For more recent discussions of the role of heuristics in reasoning, see HEURISTICS: THE FOUNDATIONS OF ADAPTIVE BEHAVIOR, (Gerd Gigerenzer, Ralph Hertwig, & Thorsten Pachur eds., 2011); and Conlisk, *supra* note 57, at 676–77.

capable of evaluating actions the more uncertain they are about those actions' results: some such actions may have a high probability of causing harm while others may have a low probability of causing harm, yet under uncertainty individuals cannot assign probabilities to results and thus cannot tell which are which. Since individuals cannot distinguish among such actions, then, they cannot perform only the ones that would have a high probability of causing harm. If individuals cannot distinguish among actions based on the probability of harm, though, they could avoid those with a high probability of causing harm by simply avoiding actions about whose results they are uncertain. As Keynes puts it, "There seems, at any rate, a good deal to be said for the conclusion that, other things being equal, that course of action is preferable . . . about the results of which we have the most complete knowledge."<sup>96</sup> Such a rule would be overinclusive, applying to actions with a low probability of causing harm and those with a high one, since individuals cannot distinguish between them. But if uncertainty deprives individuals of the ability to distinguish between actions based on their probability of harm, a general aversion to uncertainty will, by guiding individuals to avoid all risks, at least guide them to avoid the excessively dangerous ones as well.

The degree of uncertainty we face, in turn, depends on the evidence we possess: the more evidence we have about the results of a particular course of action, the more confident we may be in our expectations concerning the results of acting, thereby reducing our uncertainty.<sup>97</sup> Thus, an aversion to uncertainty may be translated into a rule of conduct: we should hesitate to act when we lack evidence about the results of acting. More colloquially, one might say, the rule of conduct required of us is simply to be cautious under uncertainty. This rule is quite different from the rule proposed by interpreters of recklessness as intentional risk-creation—namely, that we

---

<sup>96</sup> KEYNES, *supra* note 15, at 360.

<sup>97</sup> Keynes employs the term "weight" to designate the amount of evidence that bears upon a particular claim. *See generally id.* at 78–86. An argument's weight, then, is distinct from the probability that it assigns to some claim: an equal number of heads and tails among a large number of flips of the same coin is a very weighty argument, but it does not assign a very high probability to the chance of heads. *Id.* at 81. Increasing the weight of the evidence does, however, often increase the degree of justifiable confidence in the conclusion that evidence supports, whatever it may be. *Id.* at 83. Because greater evidence about the results of a particular action can decrease our uncertainty concerning its results, then, the weight of the evidence counts in favor of a particular course of action: "A high weight . . . increase[s] *pro tanto* the desirability of the action to which [it] refer[s] . . ." *Id.* at 360. For contemporary discussions of the role of weight in probability theory, see Bradley, *supra* note 59, at 2.3; and James M. Joyce, *How Probabilities Reflect Evidence*, 19 PHIL. PERSP. 153, 158–62, 165–67 (2005).

must not perform actions whose harms exceed their benefits, each discounted by their likelihood of occurring.<sup>98</sup> The latter is a highly complex formula that identifies permissible and prohibited conduct based on its fine details; the former is a blunt imperative that one would—indeed, that most do—use to instruct children. But of the two, it is the former, not the latter, that is more similar to most other criminal prohibitions: the bulk of the rules of conduct enforced by criminal law—Don't hit people! Don't take their things! Don't start fires!—are blunt imperatives prohibiting broad categories of conduct, rules that in fact are employed to teach children appropriate behavior. Most such rules, of course, do not concern uncertainty; instead, they identify specific actions that individuals should not perform. These rules guide the choices of individuals who possess sufficient knowledge to recognize an action they might perform as falling within a category prohibited by law. But because we do not always know the results of our actions, some of the crude rules we are taught must guide how we act when we are uncertain. The rule to be cautious when we lack evidence about the results of our actions is one such rule. In my view, it is this rule that prohibitions on recklessness enforce.

The aversion to uncertainty that this rule requires of us is a familiar feature of human deliberation. Empirically, people do display a preference for choices with certain rather than uncertain outcomes.<sup>99</sup> And many contemporary decision theorists have defended the idea that the uncertainty associated with an option may rationally influence an agent's decision concerning whether to perform it;<sup>100</sup> indeed, some models even explicitly require agents to prefer actions concerning whose outcomes more evidence is available.<sup>101</sup> But the general principle that agents should avoid actions about whose results they are uncertain does not, on its own, state a plausible rule of conduct that individuals could employ in deliberation to identify which actions they must avoid. Criminal law enacts prohibitions: by defining an action as an offense, the law does not merely identify it as in some way undesirable but rather

---

<sup>98</sup> *E.g.*, ALEXANDER, FERZAN, & MORSE, *supra* note 8, at 263.

<sup>99</sup> For some well-known empirical examples of decisions in which individuals display risk-aversion, see Allais, *supra* note 54; and Ellsberg, *supra* note 54.

<sup>100</sup> For a recent formal model of rational choice according to which an agent's attitudes towards uncertainty rationally influence her choices, see LARA BUCHAK, RISK AND RATIONALITY 48–81 (2013).

<sup>101</sup> *E.g.*, Peter Gärdenfors & Nils-Eric Sahlin, *Unreliable Probabilities, Risk Taking, and Decision Making*, 53 SYNTHÈSE 361 (1982).

requires individuals not to perform it. And, obviously, individuals cannot follow a rule that forbids performing any action whose results are uncertain. Because certainty is so rare, and because uncertainty is not the only consideration relevant in decision-making, individuals sometimes quite reasonably perform actions despite some degree of uncertainty concerning their results. Though uncertainty rationally ought to influence individuals' deliberation, any rule to be applied in deliberation cannot simply treat uncertainty as automatically ruling out action but instead must identify when it requires individuals not to act and when they may act despite uncertainty. The requirement to act cautiously is correct but incomplete: to guide conduct effectively a criminal prohibition on incautious action must specify when individuals are legally required to be cautious, and how cautious they are required to be on those occasions.

These questions might be answered, of course, via some sort of formal model of rational choice that explicitly specifies how the uncertainty associated with an action should influence an agent's decision whether to perform it.<sup>102</sup> But though such models may correctly characterize the choices that a rational agent would make, they cannot plausibly identify rules of decision that agents could employ in deliberation to evaluate particular actions. Individuals are no likelier to reliably quantify the extent of their uncertainty than they are to reliably quantify the degree of particular risks, and formulae that determine the permissibility of actions based on a quantitative measurement of uncertainty will likely be no easier to apply than formulae based instead on quantitative measurements of risk. Thus, just as ordinary individuals cannot plausibly choose actions based on precise quantifications of the probability of harm, they also cannot plausibly evaluate actions by precisely quantifying the degree of uncertainty associated with their results. As I have noted, criminal prohibitions generally guide conduct not by requiring individuals to perform complex evaluations of each option they face but rather by promulgating simple rules that can easily be applied to identify which actions may not be performed. Just as "be cautious" is itself a simple rule of thumb, similarly simple rules of thumb are required to specify when individuals must be cautious and how cautious they must be. And it is those rules of thumb that individuals will actually employ in deliberation to determine whether any

---

<sup>102</sup> For one such formal model, see *id.*



particular action under consideration is consistent with the general principle of aversion to uncertainty.

A rule of thumb fashioned for individuals to use to make decisions when facing uncertainty, though, must confront a difficulty that does not arise in fashioning most such rules. As I have argued, individuals cannot compute the expected utility of most ordinary actions reliably enough to justifiably rely on those computations in ordinary decision-making. But most rules of thumb respond to that difficulty by requiring individuals to evaluate actions based on other features which they can identify reliably enough: perhaps individuals cannot identify and evaluate expected consequences reliably enough to determine which instances of hitting other people have positive or negative expected utility, but they certainly can identify which actions would constitute hitting other people reliably enough to apply a rule of thumb that forbids hitting other people. That is, since many rules of thumb instruct individuals to evaluate actions based on particular consequences, those rules rely on individuals' ability to reliably identify some results of their actions. But a rule of thumb concerning the appropriate degree of caution under uncertainty cannot guide conduct in this manner because the rule applies precisely in those contexts when individuals cannot reliably identify the results of acting at all—namely, contexts in which agents face uncertainty. That is, because such rules of thumb must instruct individuals how to act when they cannot rely on their own expectations of whether acting will cause particular harmful results, whatever guidance they provide individuals cannot be based on individuals' own expectations of the results of acting. But if individuals cannot rely on their own expectations of an action's results in evaluating it, what sort of evaluation could a rule of thumb require them to employ in deciding whether to perform that action?

This problem has a straightforward solution: if individuals cannot rely on their own judgment in evaluating actions, they should rely on someone else's judgment instead. As Keynes explained, "Knowing that our own individual judgment is worthless, we endeavor to fall back on the judgment of the rest of the world which is perhaps better informed."<sup>103</sup> To be sure, it would not be feasible to ask the rest of the

---

<sup>103</sup> Keynes, *supra* note 87, at 214.

world for advice before making each decision, but other methods exist for consulting the judgment of others. In particular, the judgments that others make concerning how to act under uncertainty are reflected in how other people actually behave when they face uncertainty. Thus, the standards of behavior that people generally respect in their own decision-making reflect society's collective judgment as to how individuals ought to act when they face uncertainty. Those standards may not be reflected in formal laws, but they are typically reflected in social norms—that is, in the rules that members of a society broadly employ to govern their own conduct and to critique each other's conduct.<sup>104</sup> Thus, in contexts when uncertainty undermines the reliability of their own judgment, individuals may rely on others' judgment instead by complying with general social norms concerning how to act in those contexts.<sup>105</sup> The rule of thumb individuals must follow, in short, is simply that they should not do things that nobody else would do. Or, to employ the more technical language of the Model Penal Code's definition of recklessness, individuals may not engage in “a gross deviation from the standard of conduct that a law-abiding person would observe in the actor's situation.”<sup>106</sup>

The category of actions that nobody else would perform, however, may encompass considerably more actions than plausibly ought to be the subject of criminal prohibitions. Criminal law in general—and thus the law of recklessness—concerns actions that threaten particular sorts of unusually severe harms to others, such as injury, property damage, or even death. But, obviously, not all social norms concern actions that risk causing these sorts of harms; social norms regulate behavior to achieve a wide range of purposes, ranging from the efficient coordination of behavior to the mutual expression of respect. It violates a social norm to walk down the street naked, or to take food out of a garbage can, or to fail to hold the door when leaving a building. Simply to require individuals to defer to social norms, then, would legally mandate compliance with all these rules—after all, the law-abiding person is likely to be the sort of person who complies with social norms generally. But though

---

<sup>104</sup> In H.L.A. Hart's analysis, the “general demand for conformity” and the “social pressure brought to bear upon those who deviate” are key to the existence of a social rule. H.L.A. HART, *THE CONCEPT OF LAW* 84 (1st ed. 1961).

<sup>105</sup> Keynes, *supra* note 87, at 214 (“[W]e endeavor to conform with the behavior of the majority or the average.”).

<sup>106</sup> MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962).

such behavior may be criticized on other grounds, walking down the street naked or failing to hold the door for strangers is surely not reckless. Thus, only a subset of social norms is plausibly relevant in regulating reckless behavior—in particular, norms that regulate dangerous activities in order to avoid harm to others. Prohibitions on recklessness, in my view, require individuals to follow only these norms, by evaluating whether particular actions are consistent with the social norms specifying acceptable ways of engaging in dangerous activities. In deliberation, then, individuals must apply this rule to identify the actions about whose results sufficient uncertainty exists that criminal law may justifiably require individuals to refrain from acting.

The requirement to defer to social norms concerning dangerous activities does not require individuals always to avoid uncertainty; instead, it identifies particular contexts in which caution is required, while leaving individuals free to choose at other times. Individuals should not be required always to exercise caution under uncertainty; instead, criminal prohibitions should forbid actions only when the uncertainty facing agents creates unusually compelling reasons not to act. Innovation is important to many valuable human activities, from entrepreneurship to pharmaceuticals to scholarship to fashion, and innovation is impossible unless individuals who do not know the actual or probable results of their actions act in novel ways rather than following social norms. In these contexts, individuals sometimes should not be cautious. Furthermore, even when innovation is less central and individuals plausibly should exercise caution, it is implausible that incaution should always be threatened with criminal sanctions—not every suboptimal action should be a crime. But the rule to defer to social norms concerning dangerous activities plausibly does require individuals to exercise caution in precisely those contexts where the costs of incautious action will generally be high enough to justify a legal prohibition.

The harms of incaution are likely to be highest when individuals are engaged in dangerous behavior, which is precisely that type of behavior that on the whole produces harm more frequently. Consequently, individuals who act incautiously when engaged in dangerous activities are much likelier to create high risks than individuals who act incautiously in other contexts. Thus, a legal requirement of adequate caution

is more easily justified concerning dangerous activities than other forms of conduct. By requiring individuals to defer only to social norms that regulate dangerous activities, my interpretation of recklessness requires individuals to heed uncertainty only in the contexts where incautious action is likeliest to cause harm. A requirement to exercise adequate caution when engaged in dangerous activities, though, obviously cannot require individuals never to perform those activities; instead, it must identify when sufficient uncertainty exists about an action's results that it may not be performed. An agent's evidence about the results of a particular action determines her uncertainty; thus, rules that guide the exercise of caution must somehow identify those actions about whose results evidence is excessively scarce. The requirement to defer to social norms can distinguish between actions performed under different degrees of uncertainty because the behavior of others itself constitutes evidence regarding the results of particular actions.

People generally follow social norms: individuals normally engage in dangerous activities in the manner that the relevant social norms require. And because those norms are widely followed, actions that comply with them are widely performed, thereby providing society in general with evidence about the results of actions that comply with them. Individuals consequently face relatively little uncertainty when performing actions that comply with social norms, since they benefit from the evidence provided by others' compliance with those norms. By contrast, precisely because individuals are under social pressure to comply with social norms, violations occur relatively rarely, and thus considerably less evidence exists concerning the results of violations. Because greater uncertainty exists concerning the results of actions that violate social norms than of actions that comply with them, compliance allows individuals to avoid actions about whose results relatively more uncertainty exists.

On my interpretation, then, prohibitions on recklessness prohibit individuals from disregarding uncertainty only when two conditions are met, which correspond to the law's requirements that risks be substantial and unjustifiable to be reckless. First, individuals must defer only to social norms regulating dangerous activities. Thus, the existence of uncertainty triggers the obligation to exercise caution only when

substantial risks are involved, which depends in turn on the nature of the activity in which a defendant is engaged. Second, individuals must avoid a particular action only if performing it would be unjustifiable, given the degree of uncertainty that exists concerning its results and any other relevant factors. Thus, my account agrees with interpretations of recklessness as intentional risk-creation in taking the first sentence of the Model Penal Code's definition to state the criterion for when, in theory, individuals must refrain from acting due to risk—namely, when risks are substantial and unjustifiable.<sup>107</sup> But those interpretations also take that sentence to provide the practical test to be employed in evaluating particular actions, requiring defendants and juries alike to determine whether particular actions are reckless by themselves evaluating whether the risks they create are substantial and unjustifiable. As I have argued, though, individuals (including individual jurors) lack the capacity to reliably identify the degree of risk particular actions create. And, in fact, the Model Penal Code's own commentaries do not suggest that risks should be evaluated in practice by directly determining whether they are substantial and unjustifiable. Instead, they indicate that the second sentence of the definition of recklessness, not the first, states the practical test to be employed in evaluating particular actions: "The Code proposes . . . that the jury be asked to measure the substantiality and unjustifiability of the risk by asking whether its disregard, given the actor's perceptions, involved a gross-deviation from the standard of conduct that a law-abiding person in the actor's situation would observe."<sup>108</sup>

My account follows this suggestion: because individuals are unlikely themselves to correctly identify substantial and unjustifiable risks, they should instead defer to society's collective judgments as to whether the risks particular actions create are substantial and unjustifiable, which are reflected in the norms that societies impose on individuals engaged in dangerous activities. Social norms are costly to develop and maintain; thus, societies will not develop rules regulating how to perform a particular activity unless the benefits achieved by regulation can justify those costs.

---

<sup>107</sup> MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962) ("A person acts recklessly with respect to a material element of an offense when he consciously disregards a substantial and unjustifiable risk that the material element exists or will result from his conduct.")

<sup>108</sup> MODEL PENAL CODE § 2.02 cmt. 3.

Thus, the existence of rules designed to reduce the danger of a particular activity indicates a collective judgment that the risks that activity would create, if left unregulated, are substantial. Those norms regulate behavior, in turn, by identifying ways of performing that activity that pose an acceptably low danger to others. Because such actions are routinely performed, comparatively little uncertainty exists concerning their results: performing such actions is justified because that degree of uncertainty is insufficient to outweigh the benefits of acting. But as actions deviate more greatly from typical behavior, the uncertainty concerning their results increases until it is no longer outweighed by the benefits of acting. Such actions, therefore, are unjustified. And the social norms imposed on individuals require them not to act when uncertainty reaches that point: individuals judge others to violate the relevant norms when they judge that an action cannot be justified given the degree of uncertainty surrounding its results. Societies shape the standards of behavior they expect individuals to respect based on collective judgments about which activities create substantial risks and which ways of performing those activities are unjustifiable given the risks they create. Individuals cannot themselves reliably evaluate the degrees of particular risks, but they can guide their choices according to collective evaluations of risk by deferring to social norms concerning how to engage in dangerous activities.

This account, then, identifies the rule that individuals must employ in deliberation to avoid excessively dangerous actions: they must avoid actions that are inconsistent with social norms identifying acceptable ways to engage in dangerous activities. While the rule to avoid actions whose expected harms outweigh their expected benefits can be applied only given beliefs about risk that individuals cannot reliably form in ordinary contexts, no similarly obscure beliefs are required to bring one's conduct in line with social norms; they identify prohibited actions in terms of their actual features and thus may be applied using only beliefs about the actual features of possible actions, which individuals can form reasonably reliably. But the use of social norms to guide individuals' conduct might seem to confront a different sort of obstacle: though the beliefs required to apply those rules may be feasible for individuals to acquire, the content of those rules may be more difficult to determine

than the content of the rules criminal prohibitions ordinarily require individuals to employ. The definitions of most criminal offenses explicitly define the actions individuals must avoid, thereby explicitly identifying the rules individuals must follow. But social norms have no authoritative source or interpreter. How, then, are individuals to identify the rules they must follow? And how can individuals fairly be punished for failing to follow rules that the law does not identify clearly?

Certainly, on my account of recklessness individuals cannot identify prohibited actions by consulting the texts of criminal statutes. But this limitation would prevent recklessness from effectively guiding conduct only if those bound by criminal law ordinarily rely on statutes to identify the rules of conduct they must follow. And it is implausible that they do. As John Gardner has argued, “the general criminal law . . . [is] rarely encountered in textual form, except by police officers, lawyers and those participating in trials.”<sup>109</sup> Even clearly written laws present a mass of rules that most individuals would generally prefer to avoid.<sup>110</sup> And, of course, many criminal statutes are far from a model of clear draftsmanship—the ordinary meaning of “malice,” say, bears little relation to the mens rea distinction separating murder from manslaughter. Instead, individuals know the content of the rules they are legally obligated to follow because those rules largely track the ordinary moral obligations that individuals learn simply by participating in society: as Gardner puts it, individuals know their legal obligations “by the adequate replication in the law of clear distinctions and significances which apply outside the law.”<sup>111</sup> But individuals must employ precisely this methodology, on my account, to identify reckless actions. The injunction to be cautious is as much a core commitment of ordinary morality as any of the other rules enforced through criminal prohibitions. Thus, individuals learn the social norms that govern dangerous activities from the social dynamics of ordinary interpersonal interaction, just as they learn the content of most rules that are enforced by criminal law. The failure of the definition of recklessness to identify which actions are prohibited by social norms regulating dangerous activities does not

---

<sup>109</sup> John Gardner, *Rationality and the Rule of Law in Offences Against the Person*, 53 CAMBRIDGE L.J. 502, 513 (1994).

<sup>110</sup> *Id.* (“Most people have better things to do than acquaint themselves with a mass of legal materials, however easy to read and understand.”).

<sup>111</sup> *Id.*

undermine its ability to guide conduct, then, because the social norms themselves, not the texts of criminal statutes, are the primary source from which individuals learn which actions criminal law prohibits.

That the definition of recklessness does not directly identify which actions are prohibited, furthermore, is simply an unavoidable fact about the statutory text, which prohibits creating substantial and unjustifiable risks without stating which risks are substantial and unjustifiable. And outside of purely subjective accounts on which the agent's judgment alone determines whether a risk is substantial and unjustifiable, any interpretation of recklessness must require individuals to follow some objective standard that is not explicitly stated by law. Indeed, the objective standard my account requires individuals to follow is considerably easier to identify than the standard proposed by interpretations of recklessness as intentional risk-creation. On my view, the objective standards individuals must learn concern acceptable ways of engaging in dangerous activities—that is, those standards govern actions. And standards governing actions may plausibly be learned informally, even without explicit rules, simply by observing others' behavior—in particular, by observing the actions people perform and the actions they criticize others for performing. By contrast, interpretations of recklessness as intentional risk-creation require individuals to determine the objective values of various outcomes, both positive and negative, in order to compute the expected utility of possible actions. And it is considerably less plausible that individuals could learn those objective standards simply by observing behavior. Individuals observe only the actions people do and do not perform and how others respond to those actions. Those limited observations would be compatible with a vast range of possible assignments of values to outcomes, though, and identifying how they constrain possible value assignments would be extremely difficult. Recklessness must depend on some objective standard that is not explicitly identified by law; because standards directly governing actions may be learned relatively easily, such standards are reasonable to require individuals to learn.

Because individuals are generally capable of learning the content of social norms, then, the law can effectively guide their conduct by requiring them to defer to those norms. But the absence of any authoritative statement or interpretation of the



content of social norms may well introduce considerably more vagueness into social norms than is present in rules that stated explicitly in legal texts. The content of social norms depends on the aggregated judgments of countless individuals, who may disagree with each other about precisely which actions constitute violations. And while legal procedures exist to authoritatively resolve such ambiguities in legal rules, no analogous procedures exist for social norms, and thus vagueness of this sort is likely to persist. That vagueness might seem to undermine the guidance that prohibitions on recklessness provide: if no clear boundary distinguishes violations of social norms from conformity, it is not clear how individuals could identify which actions are prohibited. But this vagueness would undermine the law's guidance only if it were to demand that individuals comply exactly with the vague boundary of a social rule. And the definition of recklessness does not require exact compliance: the Model Penal Code deems action reckless only if it involves a "gross deviation" from the relevant standard.<sup>112</sup> Even if the precise boundary of a social norm is vague, an action must violate any plausible reconstruction of it to deviate grossly from it. Thus, a rule prohibiting deviations from a social norm may guide conduct effectively even if no precise boundary identifies precisely which actions breach that norm, since those norms will clearly identify gross deviations even if they leave vague the exact scope of minor ones. Of course, the boundary separating gross and trivial deviations may itself be vague. But the rule individuals must apply in deliberation requires them to avoid violating the norm in any way, not merely to avoid violating it grossly. Vagueness in the boundary between gross and trivial deviations does not undermine the law's guidance, then, because that boundary does not guide how individuals evaluate actions.

A rule requiring individuals to defer to social norms regulating dangerous activities, then, can guide the choices those individuals make. Obviously, though, the definition of recklessness must guide courts in adjudicating violations as well as individuals in avoiding them. My account of recklessness can plausibly be applied reliably and accurately in adjudication, furthermore, because the procedures by which courts adjudicate recklessness are well designed to accurately identify the social

---

<sup>112</sup> MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962).

norms that bind individuals. In particular, verdicts in criminal trials are normally issued by juries made up of a defendant's peers—that is, by a representative sample of the community of which the defendant is a member. As ordinary members of the community, jurors will typically be a reliable source about the social norms governing it, since ordinary individuals must themselves both enforce and comply with those norms. Furthermore, the structure of a jury verdict mirrors the methods through which social norms are defined and promulgated. Social norms arise when members of a community collectively judge certain actions to be unacceptable. And since jury verdicts require unanimity, defendants will be found reckless only if the entire jury collectively condemns the defendant's conduct. The parallels between these processes, then, ensure that guilty verdicts are likely to be reached only when a social norm against the defendant's conduct actually does exist. The existence of a social norm depends on agreement among members of a community that certain conduct is unacceptable; thus, asking jurors for unanimous agreement concerning whether the defendant's conduct violated a social norm is likely to be a reliable test of whether that conduct actually did contravene any such norm.

### **III. Reckless Mental States**

By prohibiting actions performed with certain mental states, criminal offenses promulgate rules that require individuals not to act when they possess those mental states. The mental state the law requires defendants to possess in order to be reckless, then, determines how considerations of risk will affect individuals' deliberation over what to do. Definitions of recklessness that require defendants to possess a mental state directed towards the degree of the risk an action creates guide individuals to evaluate actions by first estimating that risk. But because individuals generally cannot estimate risk accurately, I have argued, such rules would guide conduct ineffectively. Instead, because individuals generally evaluate actions based on the features they actually possess, the rules the law promulgates for deliberation must explain how individuals should determine whether they may perform actions with particular sets of features. In particular, individuals must exercise adequate caution under uncertainty. The degree of caution required, in turn, is set by social norms governing how to perform dangerous activities. In evaluating particular actions, then,

individuals must determine first whether they constitute a type of activity sufficiently dangerous for society to regulate how it may be performed safely. If so, individuals may not act except in the manner society's norms permit. Because individuals cannot on their own accurately estimate the risks particular actions create, they must instead defer to collective judgments about acceptable risk-taking that are encapsulated in the norms that a society accepts and enforces.

If prohibitions on reckless conduct aim to guide individuals to exercise adequate caution, then the mens rea of recklessness must be defined so that the individuals who fail to exercise adequate caution are guilty of recklessness. This section will therefore consider which mental states the definition of recklessness should require defendants to possess to be reckless. Defendants perform intentional wrongdoing only given an intention directed towards the feature of an action that makes it wrongful—an intention to cause injury, say, or to take an item that is the property of another. By contrast, the mens rea of recklessness should not be satisfied by one single mental state directed towards some specific feature of an action; rather, I will argue that recklessness should be defined in terms of a general property of a defendant's overall intentional state. Prohibitions on recklessness require individuals to exercise adequate caution given their uncertainty concerning the results of acting: in particular, an individual exercises inadequate caution by breaching the social norms that govern dangerous activities. Neither of these norms, however, identify prohibited actions in terms of specific features those actions possess: whether an action is adequately cautious given the agent's uncertainty or whether it violates some social norm depend on the cumulative interactions between all of its features, taken together. Thus, whether an agent exercises adequate caution in acting depends on his collective beliefs and intentions about the nature of his action and the circumstances in which he performs it. Those mental states determine whether an individual is reckless: as Model Penal Code explicitly puts it, recklessness requires a gross deviation from the conduct of the law-abiding person "considering the nature and purpose of the actor's conduct and the circumstances known to him."<sup>113</sup> A reckless defendant need not possess a specific, distinctive mental state; rather, his mental

---

<sup>113</sup> MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962).

states taken together—whatever they may be—must violate a distinctive norm governing behavior.

In my view, furthermore, no additional mental state is required for defendants to be reckless. They need not possess any attitude towards the risk they create itself: they may be reckless without considering whether their actions are risky, or even if they fully believe them not to be risky.<sup>114</sup> Thus, I reject the claim, central to interpretations of recklessness as intentional risk-creation, that the risk an action creates does not reflect a defendant's rational agency and thus cannot justify punishing him unless he recognized it in choosing to act: because my account of recklessness reconceptualizes the rule individuals must apply in deliberation, it similarly reconceptualizes when a defendant engages in the sort of improper exercise of rational agency that justifies holding him criminally liable. If individuals must evaluate actions by applying a rule to be cautious, not a rule based on the degree of risk an action is expected to create, then beliefs about risk are not required to apply the rules correctly in deliberation. Thus, both defendants who fail to recognize risk and those who believe that no risk exists may be legitimately held culpable: they fail not to respond appropriately to their beliefs about risk but rather to exercise appropriate caution given their evidence concerning the results of acting, a failure that does not require any belief that a risk exists.

*A. The Nature and Purpose of Action and the Circumstances Known to the Actor*

Recklessness must be defined, on my view, to require individuals to exercise adequate caution under uncertainty. The degree of caution required, in turn, depends on social norms that identify acceptable methods of performing dangerous activities. Criminal law promulgates rules for deliberation by prohibiting actions that violate those rules, thereby guiding individuals instead to act in compliance with those rules. Thus, if offenses of recklessness aim to prohibit individuals from acting incautiously, recklessness must be defined so that an individual is reckless just in case in deliberation he fails to apply rules governing the degree of caution required in his

---

<sup>114</sup> For a different argument that these two types of defendants are reckless, see DUFF, *supra* note 3, at 157–73.

circumstances. In general, deliberation is a mental process. Agents begin with mental states directed towards features of the world, including perceptions, beliefs, and desires concerning a possible action and its circumstances, then, on the basis of those mental states, they form a volition and act upon it. Any rule that an individual applies in deliberation must govern this process, specifying what practical conclusion he must reach—that is, what he must decide to do—given his premises—beliefs, perceptions, and so forth—concerning the decision he faces. If a rule instructs individuals to reach a particular conclusion given particular premises, they fail to follow that rule by accepting its premises yet failing to draw the conclusion it requires. Because these mental states are involved in failing to follow the rule in deliberation, acting with those mental states—that is, accepting the premises of a rule but failing to act in accordance with the conclusion it requires—must satisfy the definition of some grade of *mens rea*.

What sort of mental states might serve as the premises for such a rule? When deliberating, which attitudes directed towards an action should individuals take to require the conclusion not to act? Obviously, substantive views concerning which harms the law ought to prevent will to some extent affect which mental states make acting impermissible: the law should prohibit acting given a particular mental state only when acting on that mental state would tend unjustifiably to cause a harm the law aims to prevent. In addition, though, certain general distinctions exist among the types of mental states that might play this role. The simplest sort of rule regulating deliberation would employ a single premise, requiring individuals not to act if they possess one specific mental state. An individual could apply that rule, then, simply by determining whether he possesses that mental state—that specific belief, intention, or purpose that on its own establishes whether acting is permitted. Since that mental state alone suffices for acting to be prohibited, according to the rule, as a premise it alone suffices to dictate what conclusion an individual must reach in deliberation. Consequently, it alone suffices for *mens rea*: since individuals must conclude from it that acting would be prohibited, by accepting it and nonetheless acting they fail to correctly apply the rule the law imposes on deliberation.

Criminal law widely employs this approach to define offenses. Many criminal offenses prohibit forms of intentional wrongdoing, actions intended to have the features in virtue of which they are likely to cause harm—for example, actions intended to cause death, or injury, or damage to property. The kinds of mens rea used to define these offenses, such as intention, knowledge, or purpose, are satisfied by a particular mental state directed towards some specific feature of the action performed—namely, by a belief, intention, or knowledge directed towards the feature of the action because of which it must be avoided.<sup>115</sup> For example, a defendant who acts knowing or intending his action to cause injury commits a crime.<sup>116</sup> Because this single mental state suffices on its own to make acting prohibited, offense definitions of this sort communicate a simple rule for individuals to employ in deliberation: individuals may not perform an action if they possess a mental state—whether intention, purpose, or knowledge—that represents that action as having the specific feature in question. Such rules are straightforward to apply: since they identify an action as prohibited based on a belief or intention directed towards a single feature of that action, individuals may apply the rule simply by identifying whether a possible action has that one property. If, say, it is a crime to intentionally cause injury, then individuals may comply with the law simply by not acting on the intention to cause injury.

The simplest rules promulgated by criminal prohibitions for use in deliberation, then, are framed in terms of specific features that distinguish the actions individuals must avoid. Individuals apply such rules by avoiding actions known or intended to have those features, and they violate them by intending or knowing the actions they perform to have those features. Such wrongdoing is intentional precisely because of the intention to perform an action with some specific feature that the law defines as making action wrongful. Recklessness, however, is a form of unintentional

---

<sup>115</sup> Under the common law, offenses of intentional wrongdoing were ordinarily defined to require a mens rea of intention. *E.g.*, JOSHUA DRESSLER, UNDERSTANDING CRIMINAL LAW 117–18 (8th ed. 2018). The Model Penal Code disambiguated intention into two kinds of mens rea, purpose and knowledge. *See* MODEL PENAL CODE § 2.02(2)(a)–(b) (AM. LAW INST. 1962); DRESSLER, *supra*, at 135–37. But this distinction in fact plays little role in defining offenses. *See, e.g.*, Simons, *supra* note 11, at 470–71 (“[T]he concept of purpose is surprisingly unimportant: although the Code distinguishes between purpose and knowledge in the definitional section, it only rarely distinguishes between them in the sections specifying requirements for individual offenses.”). Thus, I treat intention, purpose, and knowledge as largely interchangeable in how they guide individual conduct.

<sup>116</sup> MODEL PENAL CODE § 211.1(1)(a).

wrongdoing: defendants are guilty only of recklessness rather than of intentional wrongdoing when they do not intend the action they perform to have some feature specified by law as wrongful. And because recklessness is a form of unintentional wrongdoing, the mental state required for defendants to be reckless must be something besides an attitude directed towards some specific feature of that sort. Interpretations of recklessness as intentional risk-creation propose that the intention to create a risk that an action will have some harmful feature might replace the intention to perform an action that actually has that feature.<sup>117</sup> On this approach, then, prohibitions on recklessness would still guide individuals' deliberation through a rule that prohibits acting given a particular mental state directed towards a specific feature of actions; the difference between intentional and unintentional wrongdoing would lie only in whether that mental state represents the action as having that specific feature or merely as risking that specific feature. But the definition of recklessness would still employ a particular intention or belief directed towards a single normatively significant aspect of the defendant's action—namely, its risk of causing harm—in order to identify when acting is impermissible.

Deliberative rules framed in terms of a mental state directed towards a specific feature of actions present certain advantages: it is more straightforward for individuals to employ a rule that may be applied based only on a belief or intention directed towards one specific feature of an action or towards relatively few such features. Rules requiring individuals to evaluate actions based on the degree of the risk they create preserve this advantage, since individuals may apply those rules based only on a single belief or intention directed towards a specific feature of actions—namely, the risk they create. But, as I have argued, individuals cannot apply this rule accurately because they generally cannot accurately estimate the degree of risk that different actions create. And while all risky actions are alike in creating risk, the underlying features of those actions that are responsible for the risks they create differ widely. Because risky actions do not all share any single feature beyond being risky, then, no single rule could prohibit all such actions by identifying some belief or

---

<sup>117</sup> See *supra* notes 28–33 and accompanying text.

intention directed towards some specific feature shared by all actions that are prohibited as excessively risky.

One response, of course, might be to enact multiple different rules that each define and prohibit a particular sort of risky action. But because there are many different ways for actions to be risky, this approach would require a considerable number of rules corresponding to the different features that create risk, and those rules would both be challenging to learn in their entirety and might well omit some ways in which actions create risk.<sup>118</sup> If the law aims to promulgate some rule governing dangerous conduct that is not tied to particular ways in which actions may be risky, then, it cannot frame that rule in terms of some specific feature shared by all excessively risky actions: the only such feature they share is the excessive risk they create, which cannot effectively be employed to guide deliberation. But if a rule guiding deliberation does not define prohibited actions in terms of a belief or intention directed towards one specific feature of those actions, how might it identify the mental states from which individuals must conclude that acting would be prohibited?

Instead, rules might guide deliberation not by instructing individuals not to act given particular beliefs or intentions whose content is a specific feature of an action but rather by instructing them not to act if a particular property characterizes their mental states as a whole. On my account, prohibitions on recklessness influence deliberation not by guiding individuals' actions given their beliefs about an action's risks but rather by guiding their actions given their uncertainty about the results of acting. And individuals are not uncertain because they possess some belief or intention directed towards uncertainty itself; rather, uncertainty arises when an individual's evidence collectively fails to identify the likelihood that acting will have particular results.<sup>119</sup> On this interpretation, prohibitions on recklessness promulgate a rule that does not identify prohibited actions in terms of a belief or intention directed towards a specific feature of that action; rather, it instructs individuals not to act when a particular property—uncertainty—characterizes their mental state as a whole.

---

<sup>118</sup> Traffic laws plausibly regulate behavior in this way by specifically defining and prohibiting each of the many particular ways in which driving may be excessively risky. *See, e.g.*, N.Y. VEH. & TRAF. LAW §§ 1100–1131.

<sup>119</sup> *E.g.*, KNIGHT, *supra* note 15, at 233.



Of course, only sometimes must individuals exercise caution because of uncertainty concerning the results of acting. Some activities cause harm quite rarely; consequently, insufficient uncertainty exists concerning the results of any single instance of that activity for caution to be required. By contrast, caution is most important when activities are dangerous: since dangerous activities cause harm more frequently, greater uncertainty exists concerning the results of particular instances of those activities. Consequently, individuals must exercise caution when engaged in those activities by acting only in compliance with the social norms governing how to perform them safely. This rule identifies when caution is required because of uncertainty by identifying actions whose results are sufficiently uncertain that performing them would be insufficiently cautious—namely, those actions that violate social norms governing dangerous activities. Like the general requirement to exercise adequate caution given uncertainty, this more precise rule does not identify any one specific feature that distinguishes all reckless actions, as they might all be distinguished by creating a risk of a particular degree of causing certain harms: for an action to violate the social norms governing dangerous activities is not a specific result that it causes or specific aspect of the circumstances in which it is performed. Thus, the rule does not employ a mental state directed towards some specific feature of an action to identify when uncertainty concerning that action's results requires agents to exercise caution. Instead, uncertainty exists concerning those actions' results precisely because they violate the social norms governing dangerous activities. And since there are many different ways in which an action can violate those social norms, individuals with many different sorts of beliefs or intentions about an action can face a degree of uncertainty concerning its results that requires them to exercise caution by not performing it.

The rule that I interpret prohibitions on recklessness as promulgating, then, identifies when individuals must not act in a different manner than prohibitions on intentional wrongdoing do. Those rules specify single premises—a particular belief or intention directed towards a specific feature of actions—from which individuals must conclude that the action is prohibited. By contrast, the rules promulgated by prohibitions on recklessness do not specify any particular belief or intention directed

towards a specific feature of actions that on its own determines whether acting is permissible. Instead, those rules only indirectly identify the mental states from which individuals must conclude that they may not act: they instruct individuals not to act when their mental states collectively satisfy some condition, but many different specific collections of mental states are together capable of satisfying that condition. Individuals must exercise adequate caution under uncertainty, but no particular mental state is required for uncertainty to require caution. Instead, uncertainty depends on an individual's evidence, taken together, and many different collections of evidence could produce uncertainty of that degree.

In particular, that uncertainty exists concerning the results of actions that violate social norms governing dangerous activities. Because the criminal law's definition of recklessness does not itself state those norms—rather, it merely incorporates them by reference through its invocation of the standard of conduct observed by the law-abiding person—it does not directly identify which mental states suffice for an individual to be reckless. Instead, it states a condition that an individual's mental state must satisfy as a whole to be reckless—a condition that individuals may satisfy by possessing many different kinds of particular mental states. The rule promulgated by recklessness prohibits acting given any beliefs ascribing any combination of features to an action in virtue of which it violates a relevant social norm. An individual violates the rule, then, by acting with any set of mental states that collectively satisfy this condition.

But while the definition of recklessness identifies when defendants are reckless in terms of some general condition, not in terms of a single mental state directed towards one specific feature of actions, any reckless defendant will satisfy that condition in virtue of the particular beliefs or intentions he actually possesses concerning specific features of his action. All intentional conduct involves beliefs and intentions about the specific action actually performed. Those beliefs and intentions constitute an individual's evidence about an action's results and therefore determine his uncertainty. And individuals evaluate actions in deliberation by considering whether, given their actual beliefs or intentions about the features an action will have, their uncertainty concerning its results requires them to exercise caution by not

performing it. Social norms governing dangerous activities, in turn, specify when uncertainty requires caution. Those norms identify which actions individuals may not perform, and individuals apply them in deliberation by determining whether the action they believe or intend that they will perform breaches any such norm. An individual violates that norm, then, by acting while aware of the specific features of his action in virtue of which it constitutes a violation. And he thereby exercises inadequate caution: given his evidence about that action's results—namely, his actual beliefs and intentions about it—he faces a degree of uncertainty that requires caution. The social norm forbids such actions precisely because of that uncertainty. In deliberation, then, individuals must conclude that acting is forbidden if they believe or intend that their action will have specific features in virtue of which it violates some social norm, since those beliefs and intentions together give rise to a degree of uncertainty requiring caution. And an individual violates that rule if he possesses beliefs and intentions of that sort and nonetheless acts.

Criminal prohibitions enforce rules by defining offenses so that individuals who accept a rule's premises but do not act according to its conclusion commit an offense. Whether a defendant is reckless, then, depends on his mental states concerning the specific features of his action—that is, on the ordinary beliefs and intentions that guide his conduct. What determines whether those mental states suffice for recklessness, of course, is not whether they include a mental state directed towards some specific feature that all reckless actions must possess. Rather, the law defines recklessness by identifying a general property that must characterize the global mental state of any reckless defendant, whatever specific mental states he possesses. Thus, a particular defendant satisfies the definition of recklessness not because of any single feature towards which he possesses some mental state but rather because the mental states with which he acts globally possess some property.

Indeed, the Model Penal Code's own statement of the test defendants must satisfy to be reckless explicitly focuses on the defendant's global mental state rather than on any single mental state directed towards one specific feature of his conduct. The risk he disregards, it explains, "must be of such a nature and degree that, considering the nature and purpose of the actor's conduct and the circumstances

known to him, its disregard involves a gross deviation from the standard of conduct that a law-abiding person would observe in the actor's situation."<sup>120</sup> Whether he is reckless depends on the "nature and purpose of the [his] conduct and the circumstances known to him"—that is, on all of his mental states respecting the action he performed.<sup>121</sup> But, obviously, if any particular defendant's mental state collectively satisfies the definition of recklessness, it must do so because of certain specific beliefs and intentions that he possesses. Each particular reckless defendant, then, is reckless because of his specific beliefs and intentions—because he intentionally fired a gun in a crowded public place, or because he intentionally lit a bonfire on the floor of his apartment's living room, or because he intentionally opened a gas valve in a residential building. To be sure, different reckless defendants possess different particular mental states, but in each case those mental states—the defendant's ordinary beliefs and intentions—suffice on their own for recklessness. No further mental state is required: a defendant is reckless not because he possesses some additional distinctive mental state characteristic of recklessness but because, given the particular mental states he did possess—whatever they may be—by acting he violated the standard of conduct of the law-abiding person.

Traditionally, the common law divided crimes into specific-intent and general-intent offenses.<sup>122</sup> The basis for that distinction, however, has been extremely obscure—a leading treatise, for example, essentially disclaims any attempt to define it, instead concluding that "[t]here is no way, with confidence, to know what these terms mean."<sup>123</sup> But my account of recklessness, which is typically understood as a kind of general-intent offense, identifies a natural and intuitive basis for understanding this distinction. Many forms of wrongdoing involve mental states directed towards some specific feature of the action performed—the intention to

---

<sup>120</sup> MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962).

<sup>121</sup> The "purpose of the actor's conduct" and "the circumstances known to him," explicitly refer to mental states of the agent—namely, his purpose and his knowledge. ("Purpose" is a kind of mental state in the Model Penal Code, though perhaps in other contexts it might be understood in a different way. *See id.* § 1.13 (defining "purposely" and equivalent terms by reference to section 2.02(2)(a) of the Model Penal Code); *id.* § 2.02(2)(a) (defining "purposely" in terms of an agent's mental states).) The reference to the "nature . . . of the actor's conduct" is more ambiguous: conceivably, "conduct" (and therefore its "nature") might be understood to refer only to an action's non-mental elements. Nonetheless, the code itself makes clear that this interpretation is incorrect, defining "conduct" as "an action or omission *and its accompanying state of mind.*" *Id.* § 1.13(5) (emphasis added).

<sup>122</sup> *See, e.g.,* DRESSLER, *supra* note 115, at 132–33.

<sup>123</sup> *Id.* at 132.

cause injury, or death, or to take property belonging to another—that defendants must possess to be convicted. The mental state required to commit these offenses, then, is plausibly characterized as a specific intention, because it must be directed towards some specific feature of the defendant’s action. By contrast, on my analysis the law describes the mental state of reckless defendants in a different way, by identifying some property that must characterize their global mental state. But individuals may possess very different beliefs and intentions even if their global mental states are all characterized by some general property: because there are many ways to violate social norms governing dangerous activities, no particular mental state directed towards the same specific feature of an action must be present whenever a defendant acts recklessly, as a particular intention directed towards the same feature of an action distinguishes all instances of some kind of intentional wrongdoing. In this sense, then, offenses of recklessness require a general rather a specific intent: the law identifies the mental states required for recklessness by specifying that any set of beliefs and intentions may be reckless so long as they possess some general property as a whole, not by specifying what their specific object must be.

To satisfy the definition of recklessness, then, a defendant need not possess any particular mental state, as is required to satisfy definitions of other grades of mens rea. Instead, whether a defendant acts recklessly depends on applying social norms to the ordinary intentions with which he acts: if the action he intended to perform is inconsistent with those norms, then his action is reckless. But all such norms identify prohibited actions under a particular description: when we learn social norms from others, they describe the particular features of actions in virtue of which they constitute violations. We learn, say, not to light fires indoors, or not to speed through crowds of people, or not to open gas valves in residential buildings. When applying those norms in deliberation, then, we apply them by determining whether actions have the features that we have learned distinguish violations—whether an action involves lighting a fire indoors, or speeding through crowds of people, or leaving a gas valve open. We apply such norms correctly by excluding an action from deliberation based on the fact that it possesses such a feature, and we violate them when we recognize that an action possesses that feature yet perform it anyway.

It is possible, however, simply to be mistaken about the nature of an action that one will perform. Individuals may perform actions that violate social norms without recognizing that those actions have the features in virtue of which they violate those norms—say, accidentally opening a gas valve by misunderstanding the nature of the mechanism, or accidentally starting a fire while cooking by failing to realize that a particular item had begun to burn. Applying a particular social norm in deliberation will identify a possible action as a violation of that norm only if an individual recognizes that it has the features that the norm employs to describe prohibited conduct. Thus, if an individual acts without realizing that his action has such a feature—that it would leave open a gas valve or start a fire—then he would not conclude that the action is forbidden by applying the norm in deliberation, since he would not recognize the action under the description employed by that norm. Of course, multiple social norms may apply to a single dangerous activity, and someone who fails to recognize the action under the description employed by one norm may nonetheless recognize it under the description employed by another. An amateur who fails to realize that his attempts at stove repair will leave a gas valve open may nonetheless recognize that his action involves manipulating gas valves in ignorance of their function, and certainly a social norm might regulate conduct under that description. But whether each particular social norm excludes an action in deliberation depends on what information about that action is available to the agent who applies that norm—in particular, on exactly what intentions and beliefs the agent had about the action he performed.

Consequently, evaluations of whether a defendant was reckless cannot depend only on facts about the results or circumstances of a defendant's action; instead, a court must evaluate whether a defendant complied with the relevant social norms by considering his mental states in acting. As the Model Penal Code says explicitly, action is reckless if it grossly deviates from the standard of conduct of the law-abiding person “considering the nature and purpose of the actor's conduct and the circumstances known to him.”<sup>124</sup> Just as with forms of intentional wrongdoing, then,

---

<sup>124</sup> MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962). Two factors mentioned in this provision, the “purpose of the actor's conduct” and “the circumstances known to him,” explicitly refer to mental states of the agent—namely, his purpose and his knowledge. (“Purpose” is a kind of mental state in the Model Penal Code, though

factual ignorance about the nature of an action may exculpate by undermining mens rea. Since forms of intentional mens rea require defendants to believe or intend that their action possesses some specific feature, those types of mens rea can be undermined only by ignorance of the fact that an action possesses that specific feature. By contrast, since on my account beliefs and intentions concerning many different specific features of actions can together ground a defendant's recklessness, ignorance concerning many different kinds of facts can undermine mens rea if it concerns the features of an action in virtue of which it violated a social norm.

Prohibitions on recklessness require individuals to deliberate by considering whether their action violates a social norm governing the degree of caution required when an individual is engaged in dangerous activities. But if an individual is mistaken about the features of some action in virtue of which it would violate such a norm, then the conclusion reached by applying all relevant norms to the action he takes himself to be performing will be that acting is permissible. Such an individual acts dangerously not because he fails to correctly apply a rule in deliberation but rather because the premises to which he applies that rule are mistaken. A defendant's factual ignorance may be culpable: in particular, the mens rea of negligence concerns the failure to perceive risk, and a defendant may negligently be factually ignorant about his action.<sup>125</sup> But, in my view, prohibitions on recklessness promulgate rules governing the appropriate exercise of caution under uncertainty, not rules governing the factual inquiry that must precede action. Thus, even if a defendant does employ mistaken factual premises about possible actions in applying social norms governing dangerous activities, he is not reckless because he does not violate those norms: applying them in his deliberation would not have yielded the conclusion not to perform his action.

---

perhaps in other contexts it might be understood in a different way. *See id.* § 1.13 (defining “purposely” and equivalent terms by reference to section 2.02(2)(a) of the Model Penal Code); *id.* § 2.02(2)(a) (defining “purposely” in terms of an agent’s mental states). The reference to the “nature . . . of the actor’s conduct” is more ambiguous: conceivably, “conduct” (and therefore its “nature”) might be understood to refer only to an action’s non-mental elements. Nonetheless, the code itself makes clear that this interpretation is incorrect, defining “conduct” as “an action or omission *and its accompanying state of mind.*” *Id.* § 1.13(5) (emphasis added).

<sup>125</sup> *Id.* § 2.02(2)(c) (“A person acts negligently . . . when he should be aware of a substantial and unjustifiable risk . . .”). In chapter 3 of this dissertation, I defend an account of negligence on which it involves a type of unreasonable factual ignorance—in particular, factual ignorance resulting from a culpable failure to inquire before acting.

### B. *Social Norms and Normative Ignorance*

To be reckless, a defendant's conduct must have violated the standard of conduct of the law-abiding person, given his beliefs or intentions about the nature of his action and its circumstances. In my view, only these mental states are required for recklessness; in particular, individuals need not possess any explicit belief or intention concerning risk, uncertainty, caution, or danger. Thus, a defendant's failure to recognize those features of an action is not exculpatory. In this respect, my account differs from interpretations of recklessness as intentional risk-creation, according to which individuals cannot be held culpable for the risks their actions create unless they believe or intend those risks to exist. Defenses of this claim typically rest on an argument connecting culpability with the capacity to control one's behavior. Individuals may be held culpable only for failing to correctly exercise rational agency in controlling their conduct. The capacity to control one's conduct, however, depends on the scope of one's conscious awareness: absent conscious awareness of a reason, an agent cannot employ it to guide his exercise of rational agency. Thus, if an individual does not recognize the risk that an action creates, he cannot treat that risk as a reason governing his conduct, and he cannot fairly be held culpable for incorrectly exercising rational agency by acting despite that risk.<sup>126</sup> Since such individuals cannot fairly be held culpable, the definition of recklessness must exclude them by requiring defendants to recognize the risk their action creates in order to be

---

<sup>126</sup> For a variety of presentations of this argument, not all of which endorse its conclusion, see ALEXANDER, FERZAN, & MORSE, *supra* note 8, at 79 (“[O]ne is culpable only for acts over which one has control. If one is unaware that, say, someone has replaced the sugar on the table with poison, then one is not culpable for placing that poison in another’s coffee and thereby killing her. For although one is in control of the conduct of placing the white substance in the coffee, the mistaken belief that it is sugar deprives one of the kind of control necessary for culpability.”); DUFF, *supra* note 3, at 152–55 (describing the connections between agency and choice and between choice and belief as the standard justification for requiring defendants to consciously believe that a risk exists in order to be reckless); STARK, *supra* note 6, at 141–46 (suggesting that arguments for why recklessness should require a belief about risk typically rely on the importance of belief to choice and the importance of choice in realizing an agent’s evaluative attitudes in his actions); Ferzan, *supra* note 8, at 641 (“What we are responsible for are those actions that we can do something about—where we can decide whether to act. And to have this sort of control, we must be aware, in the introspective sense, of what we are doing.”); Finkelstein, *supra* note 31, at 593 (“[A]n agent can be thought of as performing an action for a reason under any description under which it was foreseen by him. An action is intentional under any description under which it is foreseen, and hence it becomes plausible for us to say that an agent is responsible for his action under any description under which it was done intentionally.”); and Husak, *supra* note 7, at 210 (“[T]he beliefs a person is unable to access through introspection cannot be employed in practical reasoning . . .”).



reckless for creating it.<sup>127</sup> Someone who performs a risky action without recognizing its risk would not be reckless.

This argument considers only one particular way for individuals to control their conduct—namely, evaluating actions based on the risks that they create. Its conclusion may be plausible concerning the capacity to control behavior on that basis: perhaps individuals cannot treat an action’s likelihood of causing harm as a reason in deliberation if they do not recognize it to have that likelihood of causing harm. But in my view individuals do not avoid risk by explicitly reasoning about it; instead, we evaluate whether actions are too risky by learning and applying social norms that identify when we must exercise caution because of uncertainty. On this approach, individuals must evaluate actions not based on their likelihood of causing harm but rather based on the features identified in those norms. Consequently, the capacity to control one’s conduct does not rely on the recognition of an action’s likelihood of causing harm. So long as individuals recognize the underlying features of the actions they perform, they are capable of controlling their conduct in light of those features. To be sure, individuals may still fail to respond correctly to those features: individuals may simply fail to apply the appropriate rule, given their beliefs and intentions about an action, when evaluating that action in deliberation. But this sort of failure—a failure to recognize that a particular rule applies in deliberation, not a failure to recognize the facts about an action in virtue of which the rule applies—does not undermine culpability, for unlike factual ignorance it does not undermine an individual’s capacity to control his behavior appropriately. Rather, just as ignorance of law is no defense so long as a defendant recognizes the facts about his action in

---

<sup>127</sup> ALEXANDER, FERZAN, AND MORSE, *supra* note 8 at 328 (providing a model criminal code on which a defendant’s recklessness depends on “the risk that the actor perceives that his conduct will cause a prohibited result or results.”); STARK, *supra* note 6 at 93 (“The thought is, then, that a person cannot be said to be ‘aware’ of a specific risk attendant upon a particular token of  $\phi$ -ing without believing that that specific risk exists. Linkage to the specific risk relevant to the defendant’s wrongdoing is required because it is to that wrongdoing that the defendant, as an agent, must be tied, before the personalised condemnation of a criminal conviction is deployed legitimately.”); FERZAN, *supra* note 8 at 644 (defining recklessness to require a defendant to have “consciously recognized that her conduct was dangerous”); FINKELSTEIN, *supra* note 31 at 579 (“[F]oresight of criminal harm provides a necessary condition for criminal liability.”); HUSAK, *supra* note 7 at 216 (describing the difference between recklessness and negligence as lying in “whether defendants believe they have created a risk”); GLANVILLE WILLIAMS, *Recklessness Redefined*, 40 CAMBRIDGE L.J. 252, 256 (1981) (describing a defendant who “did not have in mind to run any risk” because he “simply forgot to consider” it as negligent, not reckless); WILLIAMS, *supra* note 8 at 82–3 (if a defendant “did not think of the danger at the moment when he acted . . . because he was completely immersed in some other matter . . . Caldwell would make him guilty of recklessness, and such a judgment would be wrong”).

virtue of which it violates the law, ignorance of the rules regulating appropriate risk-creation does not undermine culpability.<sup>128</sup>

The failure to recognize the risk that an action creates does not exculpate the agent, in my view, because that failure does not in general undermine our capacity to avoid dangerous actions. Instead, individuals are ordinarily capable of avoiding dangerous actions so long as they recognize the features of those actions in virtue of which they are dangerous. Discussions of recklessness typically focus on actions that risk harm to others, which are the primary concern of criminal law. But our capacity to avoid danger is exercised more often and thus is perhaps better illustrated by the banal, mundane contexts in which our actions risk harm to our own interests. Consider, then, an unglamorous example of those risks—the possibility of breaking one’s dishes while washing them. That risk, obviously, varies between different items: while silverware likely cannot be destroyed through any dishwashing technique available to the home cook, glassware may easily be broken if it is carelessly dropped or even placed down too forcefully. Individuals therefore must wash different items differently to avoid risking damage: a metal spoon may be handled roughly and dropped into the drying rack, but a delicate wineglass must be handled with care and gently set down. Individuals, of course, do routinely avoid risk by washing their dishes in this way. What mental states are required to exercise this sort of control over one’s conduct?

Those who take the failure to recognize a risk to be exculpatory argue that recognizing the risk an action creates is necessary to control one’s conduct in light of that risk. On this view, individuals must recognize that a particular act of handling a glass roughly would create a risk in order to treat that risk as a consideration in deliberation and to avoid it by handling the glass carefully. But it does not seem to me that individuals who carefully handle the glasses they wash ordinarily decide how to wash them based on a belief about the risk of handling them roughly. Certainly, when

---

<sup>128</sup> Duff has elsewhere suggested that the failure to appreciate the degree of risk an action created might be understood as a form of ignorance of law, not of fact, and thus no defense. Duff, *supra* note 25, at 79. But he dismisses it as an “ad hoc attempt” to avoid denying that recklessness requires awareness of risk: it employs a “diluted” notion of awareness, he argues, since it does not require a defendant to be aware of the actual degree of risk an action creates. *Id.* I have argued, however, that this notion of awareness is not an ad hoc dilution but rather follows from a principled account of how prohibitions on recklessness should guide individuals’ conduct.

I wash dishes, I make a (usually successful) effort to handle glassware carefully enough that it does not break. But ordinarily I do not consider that risk when forming volitions or acting upon them, nor do I explicitly choose how to wash dishes by considering that risk.<sup>129</sup> How, then, do I exercise control over my conduct? What is my reason for a particular act of washing a thin-walled glass carefully? Simply the fact that it is made of glass. When washing dishes, that is, I perceive that a particular item is a glass, and therefore I wash it carefully. No intermediate steps are required: my reasoning moves directly from the fact that an object is glass to the intention to wash it carefully. The rule I employ in deliberation does not take risk as a premise; instead, I reason according to the rule that glass must be washed carefully.

Of course, if a different rule guided my reasoning different intermediate steps might be required: in particular, if I reasoned by applying the rule not to create risks whose expected costs outweigh their expected benefits, then I could apply that rule only by recognizing the degree of risk an action creates. But I do not apply that rule; instead, the rule I apply tells me what to do given facts about the composition of the object being washed, not given facts about the risk of breakage. Thus, I do need to recognize that an item is made of glass to control my conduct appropriately: if I mistake a glass for plastic, I cannot treat an item's being made of glass as a reason to wash it carefully. But if risk is not the reason why I wash a particular object carefully, I do not need to recognize that risk in order to wash it carefully. What premises I must accept in order to apply a rule correctly in controlling my conduct depend on what rule I am employing; thus, if I avoid risk by applying some rule that does not itself appeal to risk, I need not have any belief about risk in order to control my conduct appropriately.

According to many commentators, however, beliefs about risk actually do control conduct in these situations. A person asked why he exercised caution—why, say, he handled a glass carefully when washing it—would almost surely answer by identifying the risk he thereby avoided—in this case, the risk of breaking the glass.<sup>130</sup>

---

<sup>129</sup> For a similar point made in the context of driving, see DUFF, *supra* note 3, at 160.

<sup>130</sup> *E.g., id.* at 61; STARK, *supra* note 6, at 98, 115; Ferzan, *supra* note 8, at 633–34; Douglas Husak, *Distraction and Negligence*, in *PRINCIPLES AND VALUES IN CRIMINAL LAW AND CRIMINAL JUSTICE: ESSAYS IN HONOR OF ANDREW ASHWORTH* 81, 88–89 (2012); Husak, *supra* note 7, at 212–13.

Indeed, even without being asked he surely could think about the risk while washing the glass, if he wanted to.<sup>131</sup> If we would cite a risk if asked why we acted cautiously, these commentators argue, that risk must have guided our conduct, but only in some non-explicit manner. Duff, for example, suggests that cautious behavior is guided by “tacit rather than explicit” awareness of risk, which “may guide [an agent’s] actions and reactions without any such conscious process of contemplating his surroundings or calling his latent knowledge to mind.”<sup>132</sup> The fact that we do not explicitly govern our decisions based on considerations of risk in deliberation shows only that those considerations must be present to our minds in some other form.

This argument interprets an agent’s answer to the question of why she acted cautiously as giving a psychological explanation of her behavior. But explanation is not the only function of the reasons we give for our actions; reasons motivate behavior but they also—perhaps more often—justify it. If asked why one performed a particular action, that is, one might describe either the psychological states that caused it or the facts that made it a correct response to its circumstances. And there is one good reason for construing an individual’s mention of risk as justifying his behavior rather than explaining it—namely, it is obviously correct as a justification. The risk that washing a glass roughly will cause it to break is the reason why people should instead wash it carefully. But even if a reason adequately justifies an action, it need not explain why the agent acted as he did. An individual asked why he washed a glass carefully, for example, might not respond by mentioning only the risk of breaking it. He might go on to explain that he sought to avoid breaking it in order to avoid the cost of replacing it, and that he sought to avoid the cost of replacing it because he is saving everything he can for retirement, and that he is saving everything he can for retirement so he can travel the world once he turns fifty. Surely, though, the fact that these reasons justify his caution does not suggest that each act of handling a glass gently must be explained by a belief about the risk that rougher

---

<sup>131</sup> DUFF, *supra* note 3, at 160 (“[W]e do indeed sometimes make our knowledge of what we are doing explicit to ourselves in such silent mental reports . . .”); STARK, *supra* note 6, at 115–16 (“Through attending consciously to the matter, Ellie can render explicit that which was guiding her conduct.”).

<sup>132</sup> DUFF, *supra* note 3, at 80; Ferzan, *supra* note 8, at 633 (“To be rational, the driver, when asked, ‘Why is this dangerous?’ must have an answer—‘This action is dangerous because it risks lives, injury, and property damage.’ Given that the actor must have a sense of what he means when he thinks that an activity is dangerous, it must figure into his practical reasoning about whether to engage in the activity.”).

handling would pose to the agent's future travel plans. And, certainly, it would be absurd to suggest that if he does not recognize the risk that roughly handling the glass poses to those travel plans, he would lack the capacity to control his conduct appropriately. He could not have handled it gently for the reason that rough handling would marginally threaten his plans to travel the world unless he recognized that reason. But he can handle an object gently for the reason that it is made of glass based only on the belief that it is glass; no further belief about how handling it roughly would implicate the deeper reasons that justify the rule is required to follow the rule.

It may seem, however, that this argument construes the explanatory role of risk too narrowly. Perhaps a particular instance of careful glass-washing cannot be explained by the belief that roughly handling that very glass on that particular occasion would risk damaging it. But the fact that an individual cites that risk to justify his decision to wash it carefully might indicate that the risk plays a different sort of role in explaining his conduct: in particular, his beliefs about risk could explain why he generally employs that rule in washing glass. No doubt this explanation is often correct: surely many people do decide to wash glass gently because of the risk that handling it roughly will break it. If beliefs about risk play only this role in explaining why individuals handle glass gently, though, then an individual's failure to perceive the risk created by a particular instance of washing a glass roughly does not undermine his capacity to control his conduct on that occasion.

Individuals do not re-derive rules for washing dishes from first principles every time they wash dishes: when acting on particular, everyday occasions we almost never decide what to do by reasoning from first principles, a cognitively taxing process liable to err. Instead, we employ more specific rules derived from those first principles that may be applied directly to possible actions, such as the rule to wash glasses gently. And once an individual has learned that rule, he can apply it in the future without referring back to the considerations that originally led him to adopt it. Indeed, some individuals may have forgotten why they first adopted the rule, and others may have adopted it without considering risk at all—perhaps, say, they were simply told to handle glass gently. These individuals would be capable of by handling glasses gently even were they unable to perceive the risk because they entirely lack

knowledge that the risk existed. And if these individuals retain the capacity to wash a glass gently even though they are unable to perceive the risk that rough handling will cause it to break, then surely an individual who knows that the risk exists in general need not perceive it on a particular occasion to control his conduct appropriately. Of course, nothing would prevent someone who does know that risk to exist from considering it on particular occasions, should he want to. But there is no reason to think that an individual who retains the capacity to recognize the general risk of roughly handling glass need employ different mental states to guide any particular instance of washing a glass than an individual who has lost or who never possessed that capacity. The perception of risk is not required for individuals to control their conduct appropriately.

Neither law nor morality ordinarily regulates how people wash their own glasses, of course. But the example illustrates more general principles about the beliefs individuals require to control their conduct: individuals can act for a particular reason so long as they recognize the existence of that reason, but they need not recognize any deeper justifications for why they should act according to that reason. How, then, does this principle apply to reckless conduct? On my view, prohibitions on recklessness are ultimately justified by the dangers of acting under uncertainty. But the principle to exercise caution under uncertainty does not guide conduct on its own, since it fails to identify the degree of caution required by the degree of uncertainty that exists on any particular occasion. Instead, I have argued, societies develop more specific norms regulating particular dangerous activities, which identify ways of performing those activities whose results are less uncertain precisely because most individuals typically act in compliance with those norms, thereby generating evidence concerning the results of compliance. The norms are justified by considerations of caution and uncertainty. But to comply with such a norm individuals need not perceive that violating it would involve inadequate caution given the uncertainty concerning the results of acting; rather, they need only perceive the features of the action in virtue of which it violates such a norm. Individuals should not open gas valves in residential buildings or light fires indoors because of the uncertainty concerning what results those actions would cause. But they need not

perceive that uncertainty or judge that it would make acting incautious in order to avoid those actions; instead, one can avoid a particular act of lighting a fire indoors, say, so long as one recognizes that it would cause a fire to be lit indoors.

Of course, even if culpability for recklessness does not require any mental state directed towards the existence of a risk, whether a defendant possesses that mental state might still affect the degree of his culpability. Under the Model Penal Code, many types of reckless conduct constitute a more serious kind of offense if the crime is performed “under circumstances manifesting extreme indifference to the value of human life.”<sup>133</sup> Such indifference, for example, can elevate manslaughter to murder,<sup>134</sup> assault to aggravated assault,<sup>135</sup> and reckless endangerment from the second to the first degree.<sup>136</sup> Different states employ different standards to determine whether a defendant manifests extreme indifference; in particular, not all understand it to involve a distinctive mental state.<sup>137</sup> When a particular mental state is required for defendants to manifest extreme indifference, though, it often must be directed towards the harm risked by reckless conduct, not merely its risky features. In order to constitute murder in California, for example, reckless homicide must be “an act, the natural consequences of which are dangerous to life, which act was deliberately

---

<sup>133</sup> MODEL PENAL CODE § 210.2(1)(b) (AM. LAW INST. 1962). Older formulations, such as Blackstone’s, speak instead of “the dictates of a wicked, depraved, and malignant heart.” 4 WILLIAM BLACKSTONE, COMMENTARIES \*199. Other state codes employ variants on or combinations of these formulas. *See, e.g.*, N.Y. PEN. CODE § 125.25 (“circumstances evincing a depraved indifference to human life”). The phrase originated as an explication of the notion of malice, which was required for a homicide to be murder rather than manslaughter; some states continue to use “malice” to describe the standard for murder, though it is substantively identical to the standards framed in terms of “extreme indifference” used for other crimes. *See, e.g.*, Commonwealth v. Packer, 168 A.3d 161, 168 (Pa. 2017).

<sup>134</sup> Compare MODEL PENAL CODE § 210.2(1)(b) (“[C]riminal homicide constitutes murder when . . . (b) it is committed recklessly under circumstances manifesting extreme indifference to the value of human life.”) with *id.* § 210.3(1) (“Criminal homicide constitutes manslaughter when: (a) it is committed recklessly . . .”).

<sup>135</sup> Compare MODEL PENAL CODE § 211.1(1) (“A person is guilty of assault if he: (a) . . . recklessly causes bodily injury to another . . .”) with *id.* § 211.1(2) (“A person is guilty of aggravated assault if he: (a) . . . causes [serious bodily] injury . . . recklessly under circumstances manifesting extreme indifference to the value of human life . . .”).

<sup>136</sup> Compare N.Y. PEN. CODE § 120.20 (“A person is guilty of reckless endangerment in the second degree when he recklessly engages in conduct which creates a substantial risk of serious physical injury to another person.”) with *id.* § 120.25 (“A person is guilty of reckless endangerment in the first degree when, under circumstances evincing a depraved indifference to human life, he recklessly engages in conduct which creates a grave risk of death to another person.”).

<sup>137</sup> Compare People v. Feingold, 852 N.E. 2d 1163, 1167 (N.Y. 2006) (“We say today explicitly [that] depraved indifference to human life is a culpable mental state.”) with Commonwealth v. Packer, 168 A.3d 161, 169 (Pa. 2017) (“One legal scholar has defined the point of demarcation for malicious conduct under Pennsylvania law as ‘dangerousness’ — ‘the . . . act creates such a dangerous situation’ that the resultant deaths or serious bodily injuries ‘are products of malice.’ (quoting Bruce A. Antkowiak, *The Art of Malice*, 60 RUTGERS L. REV. 435, 471 (2008)) (ellipsis in original)).

performed by a person who knows that his conduct endangers the life of another and who acts with conscious disregard for life.”<sup>138</sup> This test elevates a defendant’s culpability only if he recognizes the harm his action risks: murder requires knowledge that an action would endanger the life of another. But while reckless acts may be more culpable if performed in recognition of the harms they risk causing, recklessness remains culpable to some degree, albeit a lesser one, when that recognition is absent.

If individuals must regulate their conduct not by applying a rule directly concerning risk but rather by applying various rules regulating the aspects of their conduct that make it risky, no belief about risk is required to appropriately control one’s conduct. And if a defendant’s fault consists in his failure to avoid actions with some particular risky feature, not his failure to avoid risk, the failure to perceive a risk is no obstacle to his culpability. But although on my account no belief about risk is required for individuals to control their conduct and thus to be culpable for failing to control it, individuals might still require some other mental state beyond a recognition of the features of a risky action in order to control their conduct appropriately. In particular, if individuals must regulate their conduct according to social norms specifying the degree of caution to be exercised when engaged in dangerous activities, it might seem as though they must recognize the applicability of such a norm to a particular action in order to control their conduct appropriately on that occasion. For, it might seem, one cannot treat an action’s violating a norm as a reason not to perform that action without recognizing that the action would violate the norm. Thus, even if culpability for recklessness does not require a perception of risk, perhaps it should require defendants to recognize that the social norm they violate applies to their conduct.

But this argument subtly misunderstands how individuals must control their conduct so as to avoid creating risks: the demand is not to avoid actions because they constitute rule violations but rather to avoid them because of the underlying features

---

<sup>138</sup> *People v. Knoller*, 158 P. 3d 731, 738–39 (Cal. 2007) (quoting *People v. Phillips*, 414 P.2d 353, 363, (Cal. 1966)). *See also* DRESSLER, *supra* note 115, at 527 (“[I]f a parent knowingly ignores her child’s need for food or medical care to survive, the ensuing death may constitute murder. If the parent is unaware of the peril, but should be, the offense is manslaughter.”).



in virtue of which they violate those rules. Thus, individuals need not mentally apply a particular social norm in order to avoid violating it. In this respect, the sort of control individuals must exert over their conduct to avoid recklessness is identical to the sort of control they exert in complying with all sorts of criminal prohibitions. Indeed, individuals ordinarily avoid violating criminal prohibitions without thinking about those prohibitions at all: someone walking down the street, for example, can ordinarily avoid assaulting the people he passes or robbing them of their possessions without reminding himself that legal or moral rules prohibit such actions. Indeed, it would be concerning if an individual could avoid those actions only by explicitly reminding himself that such actions are prohibited.<sup>139</sup> Individuals require the capacity to recognize which actions have which features: someone who cannot distinguish assaults from other actions cannot avoid only assaults, just as someone who cannot distinguish red from green cannot interpret traffic lights, or at least not based on the different colors of the different lights. But individuals who can recognize when an action would constitute assault can choose not to perform it on the basis of that recognition alone, without explicitly considering the rule prohibiting it. So too with the social norms enforced by prohibitions on recklessness. An individual, say, stopped at a red light must comply with various rules: he may not run over a pedestrian crossing in front of him, which would be (at least) assault, and he also may not dart through the intersection in the face of oncoming traffic, which would be reckless. Individuals can—and routinely do—avoid both sorts of actions without considering the rules that prohibit them at all.

Because individuals need not recognize the applicability of a rule in order to control their conduct by acting according to that rule, criminal culpability generally does not require individuals to recognize that their action violates a rule: *ignorantia juris neminem excusat*.<sup>140</sup> Recklessness, on my account, follows this approach: an individual need not recognize that a particular action would violate a social norm to

---

<sup>139</sup> See, e.g., Bernard Williams, *Persons, Character, and Morality*, in *MORAL LUCK* 1, 18 (1981) (arguing that an agent has “one thought too many” if he is motivated not simply by features of actions that obviously are normatively significant but rather must also think explicitly that they are normatively significant).

<sup>140</sup> See 4 WILLIAM BLACKSTONE, *COMMENTARIES* \*27.

be culpable for the violation.<sup>141</sup> And recklessness, thus understood, is no different from any other kind of crime: culpability for intentional wrongdoing also does not require an individual to recognize that his action violates a criminal prohibition. Some individuals may be so distracted, preoccupied, or exhilarated by their conduct that they perform dangerous activities in violation of governing social norms without recognizing their conduct to violate the rule, firing guns in dense residential neighborhoods or lighting fires in apartment buildings without ever thinking about whether some norm prohibits those actions. But some individuals may likewise be so caught in the grip of rage that they simply do not think about whether punching someone would violate a rule against assault, or so caught in the grip of lust or drunkenness that they simply do not think about whether engaging in nonconsensual sex would violate a rule against rape. None of these failures to consider whether a rule applies undermines the defendant's culpability. The law demands that individuals control their conduct by taking certain features of actions as reasons not to perform them—the fact that an action would injure another, say, or the fact that it would create a fire inside an apartment building. An individual fails to comply with this demand so long as he fails to respond appropriately to that feature of an action, either by considering whether to treat it as a reason, then deciding not to, or by simply never considering it at all. In either case, he is culpable.

Thus, the law correctly distinguishes between the failure to recognize that an action violates a rule and the failure to recognize that features of the action in virtue of which it violates the rule. Criminal prohibitions promulgate rules for individuals to employ in deliberation by defining actions performed in violation of those rules as offenses. Those rules require individuals to conclude, based on certain mental states directed towards an action, its results, and its circumstances, that performing it would be prohibited; performing an action fails to comply with those rules when an agent

---

<sup>141</sup> To be sure, some have argued that the law should recognize a general defense of ignorance of law. *See generally* DOUGLAS N. HUSAK, *IGNORANCE OF LAW: A PHILOSOPHICAL INQUIRY* (2016). But if these arguments are right, they would not pose any distinctive objection to my theory of recklessness. The law does not currently recognize any such defense, and, consistent with that approach, my account of recklessness makes no allowance for ignorance of the rules. Perhaps the law's approach is wrong, and if ignorance of law should generally constitute a defense, it should constitute a defense to recklessness, as well. But the question of whether a defense of this sort should generally exist is simply orthogonal to the proper analysis of recklessness as a form of *mens rea*.

possesses mental states from which the rule, if applied, would entail that acting is prohibited. The failure to recognize a particular feature of an action undermines the agent's culpability, then, because in virtue of that failure he does not actually violate the rule: if, in deliberation, an agent correctly applies a rule prohibiting lighting fires indoors to what he believes would be an act of cooking dinner, he will conclude that he may perform the action, even if in fact this particular act of cooking dinner will light his kitchen on fire. Factually ignorant agents are not culpable (unless they are culpable for the factual ignorance itself) because they deliberate in compliance with the rules they are required to employ.

But an agent who fails to recognize that his action violates a particular rule—who, say, intentionally lights a fire in his living room without noting that his action violates a rule against lighting fires indoors—does not deliberate in compliance with the required rules. For one rule in particular requires him to infer from the fact that an action would involve lighting fires indoors that he may not perform it. Because he fails to draw that inference, he fails to deliberate in compliance with the rule. Thus, a defendant's failure to recognize the applicability of a rule does not undermine the claim that he failed to apply it in deliberation: he failed to deliberate in accordance with it precisely because he failed to notice that it applied. The definitions of criminal offenses identify the rules individuals must follow in deliberation by identifying actions that violate those rules, and an individual who recognizes the features of an action based on which he must decide not to perform it, but instead decides to perform it, fails to draw the practical inference he is required to draw. The failure to draw that inference constitutes a failure to follow the rule regardless of whether it is accompanied by any more or less explicit belief that failing to draw that inference would violate a rule. Consequently, to correctly identify the rule with which individuals' deliberation must comply—namely, a rule that requires them to avoid actions they recognize as having certain features—criminal offenses must be defined to require individuals only to recognize an action to have those features, not also to possess some additional mental state concerning the fact that an action would violate a rule because of those features.

On my view, then, culpability for recklessness does not require individuals to recognize that their actions violate a social norm regulating a dangerous activity for the same reason that culpability for intentional wrongdoing does not require individuals to recognize that their actions violate a law prohibiting intentionally causing a certain harm. Nonetheless, whatever structural similarities may exist between the failure to recognize the applicability of each sort of norm, differences between the two kinds of norms may suggest that treating those failures alike would be inappropriate. By defining and prohibiting criminal offenses the law requires defendants to respond appropriately to their recognition that an action possesses certain features. Perhaps this demand may be reasonable when the relevant features of the action directly cause harm, as in rules prohibiting causing death or injury: the normative significance of harm, perhaps, is so obvious that anybody who recognizes an action's harmfulness can respond appropriately. But prohibitions on recklessness, in my view, do not require individuals to respond simply to the fact that an action would cause harm; instead, such rules require individuals to respond to the features in virtue of which those actions are likely to be harmful. And because the normative significance of those features is surely less obvious, it may seem unreasonable always to hold individuals culpable for failing to accord appropriate influence to those features in evaluating actions.

This observation is surely correct to some extent. Certainly, the normative significance of harm is exceedingly obvious—perhaps nothing has a more apparent normative force. But to conclude that individuals therefore cannot be culpable for failing to respond appropriately to dangerous features of their actions seems to me to exaggerate in both directions the difference between those features and the features of actions that define intentional crimes. The normative significance of some sources of risk is sufficiently obscure that individuals cannot plausibly be expected always to respond appropriately in deliberation—it would be unreasonable, say, to expect drivers always to avoid making turns in any manner that is in fact suboptimal for tire traction, given how subtle and complex those distinctions might be. But avoiding reckless wrongdoing, on my view, does not require individuals to attend to features of their conduct that are this obscure. Expert practitioners may follow highly specialized

rules, but recklessness requires individuals to conform only with the standard of the “law-abiding person”<sup>142</sup>—that is, only to follow those very basic norms that societies apply universally, to all people. Societies enact and enforce such norms only to prohibit actions they judge to be very dangerous, to avoid risks that the Model Penal Code terms “substantial.”<sup>143</sup> Furthermore, only “gross deviations” from those norms constitute recklessness.<sup>144</sup> Thus, individuals must recognize the normative significance only of egregiously dangerous kinds of actions—pointing firearms at people, driving wildly through crowds, lighting fires in wooden buildings, and so forth. The relevance of these features of actions to an agent’s deliberation is surely less obvious than the relevance of harm. But the normative significance of highly dangerous features of actions is still extremely obvious, and demanding that individuals respond appropriately to the recognition that an action would have one such feature seems hardly less reasonable than demanding that individuals respond appropriately to the recognition that an action would cause harm.

More significantly, though, to focus on the obviousness of the normative significance of harm suggests, incorrectly, that criminal prohibitions on intentional wrongdoing expect individuals only to respond appropriately to the harmfulness of their conduct. In fact, though, most intentional offenses do not require that a defendant actually intend to cause harm at all.<sup>145</sup> Rape requires a lack of consent;<sup>146</sup> theft requires that the property taken belong to another;<sup>147</sup> fraud requires that a statement be untrue (or sometimes only misleading);<sup>148</sup> forgery requires that a document be altered.<sup>149</sup> Of course, actions with these features are prohibited because they characteristically tend to be harmful. But these features are conceptually distinct from the harm they tend to cause: these prohibitions require individuals engaged in

---

<sup>142</sup> MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962). It is perhaps noteworthy that while the standard for recklessness is defined by reference to the “law-abiding person,” the definition of negligence instead speaks of the “reasonable person.” *Id.* § 2.02(2)(d). Inasmuch as the former standard is plausibly lower—the world is all too full of unreasonable people, but most do manage to abide by the law—this difference emphasizes that only very egregious conduct constitutes recklessness.

<sup>143</sup> MODEL PENAL CODE § 2.02(2)(c).

<sup>144</sup> MODEL PENAL CODE § 2.02(2)(c).

<sup>145</sup> Indeed, a gap may exist between the intention required for guilt and the harm thereby caused even by offenses like homicide, if some killings (say, medical euthanasia for terminally ill patients) are not harmful.

<sup>146</sup> *E.g.*, State in Interest of M.T.S., 609 A.2d 1266, 1277 (N.J. 1992).

<sup>147</sup> *E.g.*, MODEL PENAL CODE § 223.2.

<sup>148</sup> *E.g.*, Securities Act of 1933 § 17(a)(2), 15 U.S.C. § 77q(a)(2) (2018); MODEL PENAL CODE § 224.7(6).

<sup>149</sup> *E.g.*, MODEL PENAL CODE 224.1(1).

particular kinds of activities to accord importance in deliberation to certain features in virtue of which acting would be likely to cause harm.<sup>150</sup> Individuals engaged in sexual activity must consider their partners' consent; individuals interacting with physical objects must consider those objects' ownership; individuals engaged in business transactions must consider the truth of their statements; individuals writing or altering documents must consider their accuracy. The failure to accord these considerations appropriate weight is culpable. Recklessness, on my view, simply extends this approach to other dangerous activities. Individuals handling firearms, or operating heavy machinery, or lighting fires must consider certain features of those actions when deliberating about how to perform them. Both sorts of rules do not directly concern harm; instead, each concerns features of actions that cause them to be risky. Just as individuals need only recognize the feature that defines an intentional offense to commit it, they need only recognize the highly dangerous features identified in a social norm to be reckless: the normative significance of the latter features, no less than the former, is obvious enough that the failure to respond appropriately is culpable.

Of course, differences do exist between the kinds of rules promulgated by prohibitions on intentional wrongdoing and prohibitions on recklessness. Unlike intentional offenses, whose definitions explicitly identify the features of actions that must regulate individuals' deliberation, I have argued that offenses of recklessness merely require individuals to comply with certain social norms. The content of those norms is plausibly less clear and more complex than the content of the rules imposed by offenses of intentional wrongdoing: while the definition of theft, say, explicitly enumerates relatively few features of actions to which individuals must respond appropriately in deliberation, social norms concerning firearm use or driving plausibly require individuals to consider to a greater variety of less precise features of their actions. Avoiding recklessness by appropriately responding to all those features is plausibly more challenging than merely attending to the explicitly defined elements of intentional crimes.

---

<sup>150</sup> For a defense of the view that consent does not directly affect whether an action is harmful but rather determines whether individuals may justifiably risk harming one other, see Renée Jorgensen Bolinger, *Moral Risk and Communicating Consent*, 47 PHIL. & PUB. AFF. 179, 186–88 (2019).

Even if the greater difficulty individuals face in complying with prohibitions on recklessness is compatible with culpability for reckless conduct, as I have argued, that difficulty should influence how culpable reckless defendants are. First, because an individual's failure to reach a higher standard in his conduct plausibly reflects a lesser shortcoming, reckless wrongdoing is less culpable than intentional wrongdoing. The law reflects this view by grading recklessness as a less serious form of culpability than purpose, knowledge, or intention.<sup>151</sup> Thus, some actions would constitute crimes if performed with purpose or knowledge but not if performed merely recklessly,<sup>152</sup> while other crimes are graded more severely if performed with purpose or knowledge than if performed recklessly.<sup>153</sup> Second, a number of doctrines, including mental defect,<sup>154</sup> immaturity,<sup>155</sup> and involuntary or pathological intoxication,<sup>156</sup> provide that individuals may sometimes be excused from criminal responsibility because of their diminished capacity. Plausibly, whether diminished capacity provides an excuse from responsibility for the failure to comply with a rule depends on the difficulty of compliance: an individual's capacities may suffice for compliance with very basic rules but not with more complex or vague ones. Thus, if complying with all the social norms governing a particular dangerous activity is more difficult than avoiding offenses defined explicitly by law, excuses for diminished capacity should apply more extensively to reckless than to intentional wrongdoing.<sup>157</sup> I will not here attempt to work out the details of how such excuses would be applied in practice. But allowing a defendant's recklessness to be evaluated in light of his diminished capacity recognizes that defendants are culpable simply for failing to comply with legal rules,

---

<sup>151</sup> The Model Penal Code, for example, explicitly provides that knowledge and purpose involve sufficient culpability to satisfy an offense definition requiring recklessness, but recklessness does not involve sufficient culpability to satisfy an offense definition requiring knowledge or purpose. MODEL PENAL CODE § 2.02(5) (AM. LAW INST. 1962).

<sup>152</sup> *E.g.*, MODEL PENAL CODE § 224.11 (defining the various possible elements of fraud in insolvency, each of which requires a mens rea of either knowledge or purpose).

<sup>153</sup> *E.g.*, MODEL PENAL CODE § 220.2 (defining the offense of causing catastrophe as a second-degree felony if performed purposely or knowingly and a third-degree felony if performed recklessly).

<sup>154</sup> *E.g.*, MODEL PENAL CODE § 4.01-.09.

<sup>155</sup> *E.g.*, MODEL PENAL CODE § 4.10.

<sup>156</sup> *E.g.*, MODEL PENAL CODE § 2.08(4).

<sup>157</sup> For example, a prominent English case limiting the scope of recklessness concerned two boys, aged eleven and twelve, who lit a building on fire by burning newspapers inside it. *See R v. G and another* [2003] UKHL 50 (appeal taken from Eng.).

while allowing the difficulty of complying with those rules to affect the extent of their culpability.

### *C. Beliefs About Risk*

A defendant's failure to think about risk is no obstacle to culpability for recklessness, since the law requires defendants to respond appropriately to certain dangerous features of their actions, not to respond appropriately to their beliefs about risk. Nonetheless, some individuals obviously do form beliefs about whether their actions will create risk. Some act recklessly having recognized the risks their actions create, but others create risks in the mistaken belief that their actions are not risky. Those who interpret recklessness as intentional risk-creation generally deny that such defendants are reckless: as Glanville Williams puts it, recklessness "may possibly be applied, by a very stern judge, to those whose minds are a blank on the subject of risk, but cannot be extended with any propriety to those who conscientiously believe, after consideration, that what they are doing is safe."<sup>158</sup> Furthermore, exculpating such defendants may seem compatible with the account of recklessness I have defended. On my view, prohibitions on recklessness require individuals to deliberate in accordance with social norms governing dangerous activities. Those norms specify the features of actions that individuals must take as reasons not to perform them. But surely whether an action would risk causing a particular harm is such a feature: in deliberation, individuals should regard the fact that an action will not cause harm as a reason why performing it would be permissible. Thus, it may seem, an agent who acts based on the belief that his action would risk no harm violates no rule in deliberation. Perhaps, then, my account should hold defendants reckless only if they both recognize the features of an action in virtue of which it is dangerous and fail to actively believe that acting would be risky.

According to this argument, the belief that no risk exists would always be exculpatory: a defendant does not consciously disregard risk if he believes no risk to

---

<sup>158</sup> Williams, *supra* note 8, at 87–88; *see also* DUFF, *supra* note 3, at 171–72 (describing how on the "choice" model of liability . . . a man who acts on an unreasonable belief . . . does not choose to take a risk"); Williams, *supra* note 127, at 282 ("[A defendant] would not be reckless . . . having ruled out any risk in his own mind . . .").



exist.<sup>159</sup> By contrast, I will argue that such beliefs only sometimes exculpate, and only if they are reasonable: an individual who acts based on the belief that no risk exists may nonetheless disregard risk recklessly. Recklessness involves a particular kind of mental state—the conscious disregard of risk. In my view, individuals consciously disregard risk when they act in the face of uncertainty about the results of their action. Whether a defendant is reckless for acting based on an unreasonable belief that no risk exists, then, depends on an analysis of the exact mental state involved in believing that an action would not be risky. In particular, whether an individual who believes no risk to exist violates rules governing appropriate behavior under uncertainty depends, at least in part, on the relationship between belief and certainty. Defendants who believe that their action poses no risk of harm might nonetheless consciously disregard risk, I will argue, because that belief is in fact compatible with uncertainty about whether an action will cause harm.

The nature of belief has been the subject of a considerable amount of recent philosophical research. Analyses of full belief often begin by noting that belief appears to involve a puzzling ambivalence towards the proposition believed. On the one hand, belief seems not to exclude the possibility of falsehood: one can accept a particular belief while recognizing that one might be mistaken. But belief also seems to involve some commitment to truth—to believe something is in some sense to accept that the world is that way.<sup>160</sup> An analysis of full belief must therefore capture how it represents a particular state of affairs but is not maximally committed to that state of affairs.<sup>161</sup> One approach analyzes this combination of attitudes by taking individuals to fully believe a proposition if they are highly confident in it: to fully believe a proposition is to take its likelihood to be greater than some particular

---

<sup>159</sup> Williams, *supra* note 8, at 90 (“Summing this up, the statement ‘I thought it was safe’ is only another way of saying ‘I was not subjectively reckless’.”).

<sup>160</sup> For a discussion of the tension between these two positions, see Sarah Moss, *Full Belief and Loose Speech*, 128 PHIL. REV. 255, 268–69 (2019); and Ralph Wedgwood, *Outright Belief*, 66 DIALECTICA 309, 316–17 (2012).

<sup>161</sup> To be sure, some have argued that despite appearances belief really does require maximum confidence: to believe something is to be maximally certain in its truth. See, e.g., Dylan Dodd, *Belief and Certainty*, 194 SYNTHESIS 4,597 (2017). But regardless of whether this view is plausible as an analysis of belief, such mental states seem irrelevant to criminal law, since individuals virtually never achieve maximum certainty in ordinary contexts.

value.<sup>162</sup> Such accounts have been widely criticized, however.<sup>163</sup> Furthermore, were belief equivalent to a degree of confidence in a proposition that exceeds a particular value, the belief that an action created no risk would be of questionable relevance to whether acting would be reckless. Different risky actions risk different harms and produce different benefits, and the values of those harms and benefits can vary widely. Whatever threshold is required to achieve full belief, then, would appear to be normatively significant only if that threshold should happen to correspond to the probability at which a result becomes so likely that the risk it will not occur is not a sufficient reason to avoid acting. And that condition often will not be met. If full belief requires estimating a proposition's probability to exceed ninety-five percent, for example, then acting with the belief that an action poses no risk of death often would be reckless, since only rarely is creating a five percent risk of death justified.

In part because this straightforward reduction of belief to confidence does not appear promising, many philosophers have instead proposed analyzing belief in terms of the role it plays in reasoning.<sup>164</sup> On this approach, to fully believe a proposition is to be willing to reason with it—to take it for granted,<sup>165</sup> or treat it as true,<sup>166</sup> or act as if it were true,<sup>167</sup> or have a policy of using it as a premise,<sup>168</sup> or frame a decision by ignoring possibilities in which it is false,<sup>169</sup> or views it as close enough to certainty.<sup>170</sup>

---

<sup>162</sup> See, e.g., Richard Foley, *Belief, Degrees of Belief, and the Lockean Thesis*, in *DEGREES OF BELIEF* 37–47 (Franz Huber & Christoph Schmidt-Petri eds., 2009).

<sup>163</sup> E.g., HENRY E. KYBURG, JR., *PROBABILITY AND THE LOGIC OF RATIONAL BELIEF* 197 (1961); Lara Buchak, *Belief, Credence, and Norms*, 169 *PHIL. STUD.* 285, 290–96 (2014); Jane Friedman, *Rational Agnosticism and Degrees of Belief*, 4 *OXFORD STUD. EPISTEMOLOGY* 57 (2013); Brian Weatherson, *Games, Beliefs and Credences*, 92 *PHIL. & PHENOMENOLOGICAL RESEARCH* 209, 211–19 (2016).

<sup>164</sup> See BRIAN WEATHERSON, *KNOWLEDGE: A HUMAN INTEREST STORY* (forthcoming); Jeremy Fantl & Matthew McGrath, *Evidence, Pragmatics, and Justification*, 111 *THE PHILOSOPHICAL REVIEW* 67–94 (2002); Keith Frankish, *Partial Belief and Flat-Out Belief* 75 (Franz Huber & Christoph Schmidt-Petri eds., 2009); Dorit Ganson, *Evidentialism and Pragmatic Constraints on Outright Belief*, 139 *PHIL. STUD.* 441 (2008); Dorit Ganson, *Great Expectations: Belief and the Case for Pragmatic Encroachment*, in *PRAGMATIC ENCROACHMENT IN EPISTEMOLOGY* 10 (Brian Kim & Matthew McGrath eds., 2019); Moss, *supra* note 160; Jacob Ross & Mark Schroeder, *Belief, Credence, and Pragmatic Encroachment*, 88 *PHIL. & PHENOMENOLOGICAL RESEARCH* 259 (2014); Brian Weatherson, *Can We Do Without Pragmatic Encroachment?*, 19 *PHIL. PERSP.* 417 (2005); Weatherson, *supra* note 163; Wedgwood, *supra* note 160. To be sure, different versions of the view differ in many important particularities, but those disagreements will not matter for my purposes here.

<sup>165</sup> WEATHERSON, *supra* note 164, at 3.4; Weatherson, *supra* note 163, at 233; Wedgwood, *supra* note 160, at 312.

<sup>166</sup> Ross & Schroeder, *supra* note 164, at 265.

<sup>167</sup> Ganson, *supra* note 164, at 454–55; Wedgwood, *supra* note 160, at 321.

<sup>168</sup> Frankish, *supra* note 164, at 85–86.

<sup>169</sup> Weatherson, *supra* note 163, at 229–30.

<sup>170</sup> Moss, *supra* note 160, at 264.

In Frank Ramsey's metaphor, beliefs are maps by which we steer.<sup>171</sup> Obviously, though, by using a particular map in a particular context we do not commit ourselves to thinking that the map represents the world perfectly. Ordinarily, a map must simplify what it represents, at least to some extent, to be usable at all.<sup>172</sup> And we choose among maps based on whether, given their simplifications and consequent deviations from reality, they are accurate enough for our purposes, given the kinds of decisions we must make. Likewise, by fully believing a proposition we do not judge that it is certainly true. Instead, we take it to be probable enough that we may rationally disregard the possibility of its falsehood. In reasoning, then, we act as if it were certain. But disregarding those possibilities in deciding how to act does not require denying that they are possible, much as using a subway map does not require that we incorrectly take its highly schematic representation of a city's layout to be accurate.<sup>173</sup>

If to fully believe a proposition is to disregard the possibilities in which it is false, though, then a greater range of considerations affect the justifiability of a particular belief than might initially appear. Certainly, whether a particular belief is justified must depend, to some extent, on what evidence supports it: the fewer the possible states of affairs in which a belief is false, the less important they are in reasoning. But whether disregarding those possibilities is justified depends also on evaluative judgments about what would result were the belief actually false: disregarding equally likely possibilities that a sandwich has been mislabeled may be justified for someone who dislikes peanut butter but not for someone with a deadly allergy.<sup>174</sup> Thus, whether individuals ought to fully believe a proposition is not simply a matter of the strength of their evidence; instead, the justifiability of full belief also depends on the importance of the interests risked by disregarding possibilities in

---

<sup>171</sup> Frank Plumpton Ramsey, *General Propositions and Causality*, in *THE FOUNDATIONS OF MATHEMATICS AND OTHER LOGICAL ESSAYS* 237, 238 (R.B. Braithwaite ed., 1931).

<sup>172</sup> See, e.g., Jorge Luis Borges, *On Exactitude in Science*, in *COLLECTED FICTIONS* 325 (Andrew Hurley tran., 1998).

<sup>173</sup> See, e.g., *Tube*, TRANSPORT FOR LONDON, <https://tfl.gov.uk/maps/track/tube>.

<sup>174</sup> Philosophical accounts in this vein often feature cases, like this one, in which the rationality of disregarding certain possibilities depends on whether trivial or serious consequences would result were they disregarded incorrectly. See, e.g., Keith DeRose, *Contextualism and Knowledge Attributions*, 52 *PHIL. & PHENOMENOLOGICAL RESEARCH* 913, 913 (1992); Ross & Schroeder, *supra* note 164, at 261.

which that belief is false.<sup>175</sup> Individuals should fully believe a proposition only in circumstances when they may effectively guide their conduct through deliberation that treats that proposition as true, which depends both on how likely the proposition is to be false and how much it matters whether it is false.<sup>176</sup> Full belief is certainty for a particular purpose; thus, individuals are justified in fully believing a proposition if the practical interests implicated by their pursuit of that purpose justify treating it as certain.<sup>177</sup>

On my account, prohibitions on recklessness regulate responses to uncertainty. Sometimes, the uncertainty individuals face concerning the results of their actions requires them not to act. If to believe a proposition were to be certain in its truth, given one's evidence, then individuals who believe there to be no risk might fall entirely outside the scope of rules regulating responses to uncertainty: an individual who acts with certainty that acting would not produce harm does not respond to uncertainty at all, it might seem, and thus his action could not constitute an impermissible response to uncertainty. Of course, were that belief unreasonable, in some sense uncertainty might still exist. Keynes and Knight each describe uncertainty as an evidential phenomenon rather than a psychological one: its existence depends on whether an agent's evidence justifies estimating risk quantitatively, not on whether an agent actually quantifies it, justified or not.<sup>178</sup> And an individual might be certain unjustifiably if he drew the wrong conclusions from his evidence; he would then face

---

<sup>175</sup> Some philosophers have argued not merely that pragmatic considerations affect whether individuals should form full beliefs but further that whether an individual actually does believe a proposition depends on whether it would be rational to disregard possibilities in which it is false. Weatherston, *supra* note 164, at 423 (“[T]he agent believes that p iff conditioning on p changes none of her preferences.”); *but see* WEATHERSON, *supra* note 164, at 3.2 (“I used to think that this could be extended to an argument that it was part of the metaphysics of belief that it was interest-relative. But . . . that isn't quite right.”). My argument will not rely on this further claim, which essentially amounts to the denial that unreasonable beliefs are beliefs at all. But it would provide an even stronger response to the view that even unreasonable beliefs can exculpate reckless defendants: if an unreasonable belief is not genuinely a belief at all, then defendants cannot be exculpated by the belief that no risk exists even though it is unreasonable, for if it would be unreasonable to hold that belief then they do not actually believe it at all.

<sup>176</sup> For versions of this point defended by different philosophers, many of whom disagree on the exact details of how practical considerations affect the rationality of belief, see WEATHERSON, *supra* note 164, at 3.3; Fantl & McGrath, *supra* note 164, at 78; Ross & Schroeder, *supra* note 164, at 274–75; and Wedgwood, *supra* note 160, at 324–26.

<sup>177</sup> To be fair, questions of full belief might sometimes implicate only theoretical interests, as when an individual is engaged in a purely theoretical inquiry that does not bear on any practical problem of how to act. *See, e.g.*, WEATHERSON, *supra* note 164, at 3.8. But however we should understand the norms governing full belief in these contexts, the criminal law is obviously concerned with beliefs that implicate more than purely theoretical interests.

<sup>178</sup> KEYNES, *supra* note 15, at 29 (doubting “whether any two probabilities are in every case even theoretically capable of comparison in terms of numbers”); KNIGHT, *supra* note 15, at 233 (describing uncertainty as “unmeasurable,” not as unmeasured).

uncertainty in what might be called an objective sense, if not in a subjective one. But an unreasonable belief on this analysis would seem to involve only a failure to properly appreciate the strength of one's evidence—namely, to overestimate the degree of confidence it justifies. And this sort of error seems not to be criminally culpable. Reasoning empirically about risk is difficult, as my arguments in this chapter have stressed.<sup>179</sup> Obviously, it would be preferable if individuals always proportioned their confidence exactly to the strength of their evidence.<sup>180</sup> But since in general nobody knows exactly how to quantify the strength of the evidence individuals possess in most ordinary cases, those individuals can hardly be faulted—much less punished—for getting it wrong themselves.

The analysis that I have outlined here, however, understands full belief to have a very different relationship with uncertainty. On that view, full belief in a proposition does not involve certainty that it is true. Rather, individuals fully believe a proposition by treating it as true when reasoning about a particular subject: they do not deny the possibility of its falsehood but rather decide to ignore those possibilities when considering a particular purpose. Full belief, then, is entirely compatible with uncertainty. Individuals fully believe a proposition not when they lack uncertainty but rather when they choose, for a particular purpose, to ignore that uncertainty in their reasoning: they recognize the existence of possibilities in which the proposition is false but they consciously disregard the risk that those possibilities obtain. Consequently, the justifiability of the decision to disregard that risk does not depend only on the strength of the evidence available to the agent; in addition, whether to consider in one's reasoning the possibility that a proposition is false depends on the practical interests implicated by that reasoning. To believe that an action creates no risk of harm, then, just is a practical decision to consciously disregard the risk that it will cause harm. And, of course, conscious disregard of risk is exactly what prohibitions on recklessness regulate. Thus, an individual does not fall outside the scope of prohibitions on recklessness by forming the belief that no risk exists. Instead, since believing that no risk exists is simply a particular way of consciously

---

<sup>179</sup> See *supra* part I.

<sup>180</sup> See, e.g., David Lewis, *A Subjectivist's Guide to Objective Chance*, in IFS: CONDITIONALS, BELIEF, DECISION, CHANCE, AND TIME 267 (William L. Harper, Robert Stalnaker, & Glenn Pearce eds., 1980).

disregarding uncertainty, individuals who act on the basis of such beliefs may act recklessly if by forming that belief they violate the rules that govern how they must respond to uncertainty.

Of course, the fact that to believe a proposition is to disregard the risk of its falsehood does not entail the absurd conclusion that individuals are reckless whenever they believe that an action will not cause harm. Uncertainty is not always a reason for caution: sometimes the conscious disregard of risk is justified, and even when unjustified its regulation through criminal law may not always be appropriate. Rather, as I have argued, prohibitions on recklessness regulate how individuals respond to uncertainty by incorporating the social norms that societies informally enforce concerning acceptable ways of engaging in dangerous activities. Those norms directly identify actions that individuals may not perform. Individuals facing uncertainty about the results of possible actions, then, should not evaluate actions based on their own estimates of the likelihood that those actions will cause harm; instead, they must evaluate actions by directly applying the relevant social norms in deliberation. Since those norms specify that individuals must avoid actions that possess certain features, individuals apply them by excluding an action from deliberation if they recognize it to have the features specified in the rule. They disregard risk recklessly, then, when their deliberation fails to comply with the rule—when they recognize that an action has those features and yet perform it anyway.

Obviously, the most straightforward way for deliberation to be inconsistent with such a rule is for an agent to violate it in a single inferential step. The rule requires individuals who recognize certain features of a possible action not to perform that action; to contravene the rule, then, one need only recognize those features of a possible action and then, on that basis, form the intention to perform it. But this is not the only pattern through which deliberation might violate the rule: in addition, deliberation might violate it if multiple inferential steps intervene between the recognition that an action has certain features and the formation of an intention to perform it. In particular, an agent might first recognize that an action has various features, including those identified in the rule. On the basis of his assessment of what features that action has, he could then decide to disregard the possibility of its causing

harm, at least for the purposes of deciding whether to perform it. Then, on the basis of that decision (and whatever other considerations might apply), he could form an intention to perform it. Deliberation of this sort, too, is inconsistent with rule: the agent recognizes the action to have the features identified in the rule, yet contrary to the rule's requirements he performs it. Because this deliberation is inconsistent with the social norms individuals are required to employ, it constitutes recklessness, too.

On the analysis of full belief I have outlined here, individuals who form the belief that an action will create no risk then perform it on that basis engage in this exact pattern of deliberation. Such deliberation is a form of practical reasoning directed towards a particular purpose—namely, evaluating whether to perform a particular action. Individuals who deliberate in this way decide to disregard certain possibilities in evaluating that action—namely, possibilities in which it would cause harm. And this mental attitude corresponds exactly to the belief that the action will not cause harm, since to believe a proposition is just to disregard possibilities in which it is false for the purpose of reasoning in a particular context. Thus, an individual who performs an action on the basis of the belief that it will not cause harm may act recklessly so long as he satisfies the same test as an individual who does not form that belief: namely, each must recognize the action to have the features in virtue of which it is prohibited by some social norm governing dangerous activities. Since each defendant would conclude, by applying that norm, that acting is prohibited, each violates that norm in his deliberation.

To be sure, the two kinds of action involve deliberation that is inconsistent with that norm in different ways. An individual who does not form the belief that no risk exists consciously disregards risk at the very same time that he decides to act: in forming the intention to perform an action recognized to have certain features he thereby violates the rule requiring him to avoid actions that have those features. By contrast, if an individual does form that belief, his decision to act is not reckless considered on its own: he decides to act because he believes that acting would create no risk, and obviously no social norm prohibits performing actions that create no risk. His disregard of risk is distinct from his decision to act: he disregards risk in forming the belief that his action will not cause harm, then he subsequently decides to act

based on that belief. But, in my view, whether a defendant is reckless need not depend solely on whether his decision to act, given the explicit considerations on which it was based, itself consciously disregarded risk, without regard to any earlier stage in deliberation. The actual formation of an intention is not the only step in deliberation that must comply with social norms concerning dangerous activities. Instead, a defendant can be reckless if his conclusion that acting would not cause harm violated those norms given the evidence on which it was based. The law, in short, can disregard the defendant's belief about risk, looking through it to the beliefs about the nature and purpose of his action and its circumstances that justified his conclusions about risk.

On my account, then, defendants may be reckless even when they believe an action to create no risk because courts may disregard those beliefs when evaluating recklessness. If my approach countenances this strategy of disregarding some of a defendant's beliefs in evaluating his recklessness, however, it might seem arbitrary for me to apply it only to a defendant's beliefs about risk. Indeed, the account of full belief I have outlined here applies to all beliefs, not only to beliefs explicitly about the risk of causing harm: whenever a defendant forms a belief about anything, he consciously disregards the risk that the proposition believed is false. And if the claim that full belief is itself a form of conscious disregard of risk justifies looking through a defendant's beliefs to the basis upon which they are formed in evaluating his recklessness, it would seem that any belief could be disregarded in this way. In particular, I have argued that a defendant must hold certain beliefs about the nature and purpose of his action and about its circumstances in order to be reckless.<sup>181</sup> But shouldn't the argument for looking through belief also apply to a defendant's beliefs about the nature and purpose of an action and about its circumstances? In forming those beliefs, too, defendants consciously disregard the risk that they have mistaken an action's nature or its circumstances. And if evaluations of recklessness may look through full beliefs to the evidence based on which those beliefs were formed, then, it might seem that a defendant could be reckless because of the evidence justifying his beliefs about an action's nature or circumstances, not only because of those beliefs

---

<sup>181</sup> See *supra* part I.A.



themselves. This process iterates: the beliefs that constitute that evidence themselves involve a conscious disregard of risk, and thus evaluations of recklessness could look through them to the basis upon which they were formed, and so forth. But this approach would expose a defendant to liability for recklessness whenever any belief fails to be justified by the evidence on which it was based. And surely this conclusion is absurd: not every error in factual reasoning is a crime, even if harm does result.

Though this conclusion is absurd, it does not follow from my account: courts cannot look through any belief a defendant holds and find him reckless by evaluating the evidence based on which he formed it. To be sure, this possibility is not precluded by my analysis of the conscious disregard of risk, since full belief does always involve consciously disregarding the risk of error. But the conscious disregard of risk, on my account, is not always reckless; instead, it is reckless only when it is inconsistent with the standard of conduct that the law-abiding person would observe, as reflected in the social norms that a society enforces concerning dangerous activities. Thus, although every full belief does involve the conscious disregard of risk, acting on that full belief can be reckless only if disregarding the risk of its falsehood violates some relevant social norm. And for the vast majority of ordinary factual beliefs that play some role in deliberation over particular actions, no social norm will govern when individuals should be willing to disregard the risk of error. Of course, often individuals will disregard that risk in error: mistaken factual reasoning is common. But it simply would not be possible for social norms to regulate every factual inference that any individual could possibly draw about any subject. We do not know the right answer to every factual question; even if we did, we could not devise a system for effectively communicating those answers to every individual; and even if we could it would not be reasonable to expect individuals to perfectly implement an explicit set of rules of such enormous complexity in their everyday decision-making. At some point, we have no choice but to rely on each other's capacities to adequately represent the world: people must be permitted to exercise

their own judgment in forming beliefs, then to act upon them.<sup>182</sup> When that judgment goes astray, as it surely will, what results is a tragedy, not a crime.

But although obviously the formation of factual beliefs in deliberation cannot always be regulated by social norms, certainly social norms can sometimes regulate how individuals should treat the risk of error in their deliberation. Evaluations of recklessness may look through a defendant's beliefs about risk to the evidence justifying those beliefs, then, not simply because to believe that acting will not cause harm is to consciously disregard the risk of causing harm but rather because on some occasions that conscious disregard violates a social norm. In most ordinary decision contexts, I have argued, individuals cannot accurately quantify the probability that their actions will cause certain results. Thus, while ideally individuals would avoid actions whose benefits do not justify their expected harms, individuals themselves ordinarily cannot identify the expected results of an action accurately enough to reliably identify which actions will, in fact, create excessive risks. Instead, individuals face uncertainty about the results of possible actions: they simply do not know what results an action will cause, and thus they cannot reliably judge whether it should be performed. In many cases, of course, individuals must simply face this uncertainty as best they can, reaching some overall judgment concerning an action's justifiability in light of the limited information they do possess about it.<sup>183</sup> But precisely because individuals' own judgment is so unreliable, they should not rely on it if an alternative exists—if, that is, they can rely on someone else's judgment instead. When social norms exist, they encapsulate that alternative judgment—a collective community judgment as to how individuals should act in the face of their own uncertainty. Because that uncertainty prevents individuals' own judgment from being reliable, they should defer to the community's judgment instead, just as they defer to the community's judgment by complying with prohibitions on intentional crimes.

But an individual who acts under uncertainty by forming and acting upon a belief that an action will not cause harm does not comply with this approach. The

---

<sup>182</sup> To be sure, the process of gathering evidence to form such beliefs may be regulated. According to the view I defend in chapter 3 of this dissertation, violating the rules regulating that process constitutes criminal negligence.

<sup>183</sup> Keynes and Knight, for example, each stress how decision-making under uncertainty involves a holistic evaluation of a possible actions rather than some sort of deduction. *See supra* notes 71–75 and accompanying text.

law, I have argued, requires individuals to defer to the community's judgments concerning how to act under uncertainty. Forming the belief that an action will not cause harm, however, is itself a method of responding to uncertainty—namely, dismissing that uncertainty by treating the absence of harm as certain in one's reasoning about what to do. And it is a method of responding to uncertainty different from what the law requires. By forming and acting upon his own belief about whether harm will result from acting, an individual relies on his own judgment about the results of his action. Doing so directly violates a norm that requires him to defer instead to a collectively judgment precisely because individuals cannot make reliable judgments given the uncertainty they face in those circumstances. To be sure, social norms do not always prohibit relying on one's own judgment when facing uncertainty; individuals may freely guide their own conduct so long as no social norm regulates appropriate behavior in their circumstances. But when such a norm does exist, an individual who instead relies on his own beliefs about risk in acting engages in the exact sort of behavior—risk-taking under uncertainty—that the norm aims to prevent.

Evaluations of recklessness cannot always look through a belief to the evidence on which it was based, I suggested, because the law must generally permit individuals to form their own beliefs and act upon them: we cannot hope to collectively regulate every factual inference that any individual might make in deliberation. But the social norms governing dangerous activities identify one particular context in which individuals are not permitted to form their own beliefs and act upon them. The norm prohibiting certain actions represents a collective judgment that their possibility of causing harm may not be disregarded, and individuals must guide their conduct according to that collective judgment, not their own individual one. The law may look through beliefs about risk in evaluating recklessness because relying on those beliefs was itself reckless: in evaluating the defendant's deliberation, it may conflate the steps of forming a belief and acting upon it, because the rule individuals must follow itself requires individuals to conflate those two steps in their deliberation. Individuals are required to conclude, from the fact that an action has certain features, that they may not perform it; thus, they violate that rule by

recognizing those features then performing the action. This description is satisfied, furthermore, regardless of whether additional deliberative steps intervene between the recognition and the action: a defendant may either recognize the features of the action and directly perform it or recognize those features, first conclude that no risk of harm exists, and only then perform it. Because the rule itself treats the intervening belief about risk as irrelevant to what action is ultimately required, defendants violate that rule regardless of whether they form that belief.

Social norms, for example, plausibly prohibit individuals from firing loaded guns at one another. (At least, I certainly understand the social norms I live under to prohibit such actions.) One way to violate that norm is to point a gun at someone and fire, believing there to be a risk of causing death (or giving no thought to that risk). But an agent also violates that norm if, before firing, he first forms the belief that there is no risk he will shoot someone—say, if he believes the gun is loaded so that the chamber will be empty when he fires.<sup>184</sup> For in each case he has fired a loaded gun at another, the precise action the norm forbids. Of course, in one case he believes that nobody will be shot, but that belief does not matter: the norm simply forbids firing loaded guns at people, at least on my understanding, rather than forbidding firing loaded guns at people except if nobody will be shot. The rule, that is, does not allow people to take the risk that they, like the defendants in *Malone* and *Lamb*, will incorrectly judge whether in a particular instance the chamber will be empty; instead, since firing loaded guns at people is of high danger and low value (at least absent some compelling justification, for which an exception might exist), the rule simply requires people not to do it. By forbidding individuals from firing loaded guns at people, it forbids individuals from disregarding the risk that such actions will cause harm. And by forming the belief that a particular such action will not cause harm, an individual consciously disregards that risk. When he acts on the basis of that belief, then, he acts recklessly.

This social norm, at least as I understand it, makes no exception for beliefs about risk at all: individuals are simply forbidden from firing guns. As I have

---

<sup>184</sup> See, e.g., *Commonwealth v. Malone*, 47 A.2d 445 (Pa. 1946); *R v. Lamb* [1967] QB 981 (quashing conviction for jury misdirection without disputing that the charged conduct could constitute manslaughter).

construed it, then, an individual could be reckless for firing a loaded gun even with a reasonable belief that nobody would be shot—for example, if there was good reason to think the chamber would be empty but a broken mechanism caused it nonetheless to be loaded when fired. This approach strikes me as reasonable for regulating this activity in particular: given the danger, people just shouldn't fire loaded guns at each other at all. But surely an individual's beliefs about whether an action will cause harm are not always irrelevant to how she should act. Landing a plane in bad weather, for example, is a dangerous activity. Unsurprisingly, then, social norms regulate how it may be performed: in general, we do not simply rely on people themselves to accurately judge when it may be done safely. But the norm regulating landing planes in bad weather cannot simply prohibit all such actions: while in many cases people certainly should not land planes in bad weather, there seem to be ways to do it safely—professional pilots do it frequently, for example—and those actions should be permitted. How, then, could the relevant norms prohibit only dangerous instances of landing planes in bad weather?

One obvious approach would be to modify the content of the relevant norm by adding an exception for safe methods of landing. That norm, then, would generally prohibit individuals from landing planes in bad weather, but it would contain an exception that articulates how to land a plane safely in bad weather. Such a norm would correctly identify the actions that individuals should and should not perform. But it might seem unreasonable to apply norms that articulate acceptable conduct with this degree of specificity in defining when conduct is reckless. I have argued that prohibitions on recklessness need not explicitly define prohibited conduct, and may instead incorporate it through reference to the standard of the law-abiding person, because societies may reasonably expect individuals to learn the content of the norms regulating dangerous activities, much like societies reasonably expect individuals to learn the content of prohibitions on intentional wrongdoing. Perhaps that claim is reasonable concerning some norms regulating dangerous activities—norms that prohibit firing loaded guns at people, or that prohibit leaving gas valves open in residential buildings, and so forth. But it is hard to see how ordinary individuals could reasonably be expected to know how to land planes safely in bad weather. A standard

that individuals must observe on pain of criminal punishment, like the law-abiding person standard, surely cannot define dangerous conduct with this degree of specificity.

Thus, the social norms that define the standard of the law-abiding person must allow an exception (at least sometimes) for safe methods of performing dangerous activities without thereby unreasonably requiring ordinary individuals to know the details of those methods. These objectives may be jointly achieved by framing the relevant norms in terms of reasonable beliefs: individuals, that is, may be prohibited from landing planes in bad weather unless they reasonably believe that doing so will not cause harm. Specialized expert knowledge is not required to learn this norm; nor is it required to comply with the norm, since even those who are utterly unable to distinguish reasonable beliefs about harm from unreasonable ones may comply simply by not landing planes in bad weather. But the exception can still permit only safe methods of landing planes in bad weather, depending on which beliefs it takes to be reasonable: in particular, whether beliefs that landing will not cause harm are reasonable may depend on whether they are formed based on reasonable principles concerning how to safely land a plane in bad weather. Individuals are thus permitted to land plans in bad weather so long as they apply those principles, since actions performed in accordance with those principles are reasonably believed not to cause harm. Consequently, a defendant may be exculpated by his reasonable belief that his action would cause no harm. But that belief exculpates only because the social norms that govern that activity permit acting based on that belief, not because the belief that an action will cause no harm in itself prevents a defendant from being reckless.

#### **IV. Conclusion**

In this chapter, I have developed an account of how individuals ought to respond to uncertainty in order to analyze how the criminal law should define recklessness. Mine is not the first account to approach recklessness from this perspective. But other theories of recklessness that focus on how prohibitions on recklessness guide conduct treat individuals as though they generally make decisions under uncertainty based on explicit beliefs about the degree of risk that actions would create. Consequently, such accounts typically argue that prohibitions on recklessness

should be defined to govern how individuals must act in light of their explicit beliefs about risk. I have argued, however, that this approach misunderstands how individuals actually do deliberate when the results of acting are unknown. Because our evidence almost never permits us to rationally identify the degree of risk an action creates, we cannot ordinarily employ beliefs about risk itself to evaluate actions. Instead, we deliberate using rules that identify prohibited actions based on their actual features. The guidance recklessness provides must involve rules of this sort if it is to be effective. But because the definition of recklessness does not directly identify which actions are too risky itself, it can only identify those actions indirectly. It does so, I have argued, by incorporating the social norms governing dangerous activities. Individuals avoid performing unreasonably risky actions primarily by learning informal rules that societies develop to govern dangerous activities. Prohibitions on recklessness enforce those informal rules through the formal apparatus of criminal law, just as by defining intentional offenses the law criminally enforces moral rules forbidding intentional wrongdoing. Whether an individual is reckless, then, does not depend on his explicit beliefs about the risks his actions create. Rather, it depends on whether he intentionally performs an action that grossly deviates from social norms governing dangerous activities—whether he intentionally performs an action that no law-abiding person would perform.

## Chapter 3:

# Criminal Negligence, Epistemic Recklessness, and the Failure to Inquire

A puzzle lies at the core of criminal liability for negligent conduct. Unlike intentional or reckless wrongdoing, negligence does not involve conscious awareness of the aspects of an action that make it criminally prohibited. In the much-cited formulation of the Model Penal Code, “A person acts negligently with respect to a material element of an offense when he should be aware of a substantial and unjustifiable risk that the material element exists or will result from his conduct,” and it is no defense that a defendant was not in fact aware of a risk that he should have perceived.<sup>1</sup> On the other hand, criminal law generally predicates liability on the defendant’s possession of some culpable mental state, all of which typically involve some sort of awareness of the features of an action that make it wrongful.<sup>2</sup> This principle is classically embodied in the Latin maxim *actus not facit reum nisi mens sit rea*,<sup>3</sup> and the intuition it reflects is widely held.<sup>4</sup> On this view, a defendant must be

---

<sup>1</sup> MODEL PENAL CODE § 2.02(2)(d) (AM. LAW INST. 1962). Similarly, negligence liability in tort is imposed for failing to act as a reasonable person would, even if a defendant is unaware of what a reasonable person would do. *Vaughan v. Menlove*, 132 E.R. 490 (1837).

<sup>2</sup> This principle distinguishes criminal law from tort law, which does not impose liability only on defendants who possess some culpable mental state, whether because tort aims at compensation rather than punishment, because its verdicts do not express the same sort of condemnation implied in criminal conviction, or for some other reason. Indeed, since prominent forms of strict liability exist in tort, it is not clear that justified tort liability even requires wrongdoing. Liability for negligence in tort law, then, escapes the puzzle that criminal negligence must confront.

<sup>3</sup> See, e.g., 3 EDWARD COKE, INSTITUTES OF THE LAWS OF ENGLAND 10 (London, E. and R. Brooke 1797). The maxim means, roughly, that no act is guilty without a guilty mind.

<sup>4</sup> See LARRY ALEXANDER, KIMBERLY KESSLER FERZAN & STEPHEN J. MORSE, CRIME AND CULPABILITY: A THEORY OF CRIMINAL LAW 70 (2009); Claire Finkelstein, *Responsibility for Unintended Consequences*, 2 OHIO ST. J. CRIM. L. 579 (2005); Jerome Hall, *Negligent Behavior Should be Excluded from Penal Liability*, 63 COLUM. L. REV. 632 (1963); J.W.C. Turner, *The Mental Element in Crimes at Common Law*, 6 CAMBRIDGE L.J. 31 (1936).



aware of the wrongful features of his action to be justifiably convicted of a crime. If both claims are true, an agent must be unaware of the features that make her action prohibited to be negligent, but she must be aware of those features to be culpable. In Holly Smith's words, then, negligence is "highly puzzling, since [it] violate[s] a standard conception of what is required for an agent to be blameworthy for her action."<sup>5</sup>

This puzzle arises from two incompatible positions: the claim that culpability requires awareness and the claim that negligence precludes awareness. Consequently, most responses to the puzzle choose one claim to accept, then develop arguments for rejecting the other. First, some accept that criminal liability requires some awareness of the wrongful features of an action, thereby rejecting criminal liability for negligence.<sup>6</sup> Individuals may legitimately be punished only for their informed, wrongful choices, and if negligence involves only an unchosen failure to appreciate risk rather than a choice to wrongfully create risk, it cannot legitimately be punished. Negligence, to be sure, is a comparatively limited form of criminal culpability, and thus eliminating liability for negligence would not work too radical a change on the criminal law.<sup>7</sup> Nonetheless, though negligence may be limited it is not insignificant, and many find it intuitively compelling that at least sometime agents who fail to appreciate risks deserve criminal punishment.<sup>8</sup> Thus, entirely foreclosing the possibility of criminal liability for negligent conduct would require at least somewhat reconceptualizing the limits of criminal liability.

While these approaches seek to revise intuitive judgments about the scope of criminal liability by rejecting criminal liability for negligence, other theorists have instead sought to revise intuitive judgments about the requirements of criminal liability. On these views, criminal liability for negligence is justified because the kind of awareness negligent agents lack is not required for criminal liability. Though fully

---

<sup>5</sup> Holly M. Smith, *Non-Tracing Cases of Culpable Ignorance*, 5 CRIM. L. & PHIL. 115–146, 117 (2011).

<sup>6</sup> See *supra* note 4.

<sup>7</sup> The Model Penal Code, for example, makes recklessness the default degree of culpability required for conduct to be criminal, MODEL PENAL CODE § 2.02(3), on the grounds that "since negligence is an exceptional basis of liability, it should be excluded as a basis unless explicitly prescribed," *id.* cmt. 5.

<sup>8</sup> For criticism of attempts to sharply distinguish negligence from higher grades of culpability, see Marcia Baron, *Negligence, Mens Rea, and What We Want the Element of Mens Rea to Provide*, 14 CRIM. L. & PHIL. 69, 76–81 (2020).

informed wrongful choice may in some sense be the paradigm of culpable wrongdoing, such views deny that it is required for culpability; instead, the same basis that grounds the culpability of wrongful choices also grounds the culpability of other sorts of conduct, including negligence.<sup>9</sup> Some middle ground exists, on these views, between innocent conduct and wrongful choice. And though these accounts differ in their details, they often analyze liability without choice in terms of the evaluative attitudes that guide a defendant's conduct—his desires, his motivations, or more generally what he cares about.<sup>10</sup> Most often, objectionable evaluative attitudes manifest themselves in the choice to do something that invades an interest the criminal law protects, which justifies liability under higher grades of mens rea. But the negligent failure to recognize a risk can itself reflect a morally blameworthy attitude of indifference to the interests of others: people who care about avoiding harm to others ordinarily notice when their actions threaten harm; and those who do not recognize when their actions may harm others simply may not care enough about whether or not they do.<sup>11</sup> Because negligent wrongdoing reflects this absence of concern, it is a legitimate basis for criminal liability.

---

<sup>9</sup> Such defenses of criminal negligence align with philosophical accounts that deny the necessity of wrongful choice to moral responsibility, and instead hold agents morally responsible for acts with different sorts of mental states in their causal histories. See, e.g., Robert Merrihew Adams, *Involuntary Sins*, 94 PHIL. REV. 3 (1985); Angela M. Smith, *Control, Responsibility, and Moral Assessment*, 138 PHIL. STUD. 367 (2008); P.F. Strawson, *Freedom and Resentment*, 48 PROC. BRITISH ACAD. 187 (1962). Although these philosophical accounts may lend support to arguments defending a broader scope for criminal liability, it is also possible (as defenders of these accounts acknowledge) that the scope of justified blame may exceed the scope of justified criminal punishment. See, e.g., Pamela Hieronymi, *The Force and Fairness of Blame*, 18 PHIL. PERSP. 115, 118 (2004); Smith, *supra*, at 378.

<sup>10</sup> R.A. DUFF, INTENTION, AGENCY, AND CRIMINAL LIABILITY: PHILOSOPHY OF ACTION AND THE CRIMINAL LAW 139–79 (1990); FINDLAY STARK, CULPABLE CARELESSNESS: RECKLESSNESS AND NEGLIGENCE IN CRIMINAL LAW 226–272 (2016); VICTOR TADROS, CRIMINAL RESPONSIBILITY 237–64 (2007); Stephen P. Garvey, *What's Wrong with Involuntary Manslaughter*, 85 TEX. L. REV. 333 (2006); Jeremy Horder, *Gross Negligence and Criminal Culpability*, 47 U. TORONTO L.J. 495 (1997); Kenneth W. Simons, *Culpability and Retributive Theory: The Problem of Criminal Negligence*, 5 J. CONTEMP. LEGAL ISSUES 365 (1994); Victor Tadros, *Recklessness and the Duty to Take Care*, in CRIMINAL LAW THEORY: DOCTRINES OF THE GENERAL PART 227 (Stephen Shute & A.P. Simester eds., 2002). Advocates of these views all agree that criminal liability may be imposed on defendants who are unaware of the risks their actions create, and thus all accept the legitimacy of liability for negligence, but some deny further that liability for recklessness should require defendants to be aware of the risk they create. E.g., DUFF, *supra*, at 141–42; Tadros, *supra*, at 257–58. And it is consequently unclear how some of these theorists would distinguish recklessness and negligence. See, e.g., STARK, *supra* note at 269–70 (noting that Tadros's account of liability without awareness of risk could equally well constitute a theory of recklessness or a theory of negligence).

<sup>11</sup> As I have noted, different accounts employ different terminology to describe the evaluative attitude that justifies criminal liability. See, e.g., DUFF, *supra* note 10, at 141 (“My action might also be condemned as displaying my disregard for, or indifference to, the harm which it causes . . .”); STARK, *supra* note 10, at 271–72 (arguing that culpable negligence must “demonstrat[e] insufficient concern for others . . . because of an accepted facet of [the defendant's] character”); TADROS, *supra* note 10, at 264 (holding culpability without awareness to be justified only for “agents whose beliefs manifest the kind of vice for which the attribution of criminal responsibility is appropriate”); Garvey, *supra* note 10, at 337 (tracing negligence to “the causal influence of a desire [the

Each of these two strategies attempts to solve the core puzzle of criminal negligence, then, by rejecting one of the two claims that give rise to it. Consequently, each must reject either the intuitively compelling view that certain apparent cases of negligence are culpable or the intuitively compelling view that culpability requires some kind of awareness of wrongdoing.<sup>12</sup> By contrast, in this chapter I will attempt to develop an account of criminal negligence that need not choose between these two claims. Because the core puzzle of criminal negligence arises from the particular analysis of negligence it presupposes, I will propose a reconceptualization of criminal negligence on which both of the two apparently incompatible claims involved in the core puzzle may be true. The standard account of negligence distinguishes it from recklessness based on the psychological attitude it involves: while recklessness involves awareness of the risk an action creates, negligence involves a risk of which the agent is unaware.<sup>13</sup> On this approach, recklessness and negligence create the same sort of risk, and negligence is distinctive because it involves a novel form of culpability for creating that risk. The core puzzle follows directly on this understanding of the distinction: if negligence is a form of culpability for risk that by definition involves no awareness of risk, awareness must not be required for culpability if negligence is culpable at all. By contrast, I will distinguish negligence and recklessness based on the kinds of risks they create. On my proposal, negligence creates a different kind of risk from recklessness, but it does not involve a distinctive kind of psychological attitude towards that risk. And if negligence involves a different kind of risk than recklessness, the apparently contradictory claims that give

---

defendant] should have controlled”); Simons, *supra* note 10, at 377–78 (“An actor who is culpably indifferent to a harmful result . . . cares much less than she should about bringing it about.”). These arguments are broadly consistent with philosophical accounts of moral responsibility that closely tie it to an agent’s evaluative or motivational states. See generally NOMY ARPALY & TIMOTHY SCHROEDER, IN PRAISE OF DESIRE (2014). It is worth noting, however, that some theorists who take criminal culpability to require indifference to the interests of others disagree that negligence reflects such indifference. ALEXANDER, FERZAN, & MORSE, *supra* note 4, at 69–81; Larry Alexander, *Insufficient Concern: A Unified Conception of Criminal Culpability*, 88 CALIF. L. REV. 931, 949–52 (2000).

<sup>12</sup> For a statement of the intuitive force of the former, see Seana Valentine Shiffrin, *The Moral Neglect of Negligence*, 3 OXFORD STUD. POL. PHIL. 197 (2017). For a statement of the intuitive force of the latter, see Gideon Rosen, *Kleinbart the Oblivious and Other Tales of Ignorance and Responsibility*, 105 J. PHIL. 591, 608–09 (2008).

<sup>13</sup> ALEXANDER, FERZAN, & MORSE, *supra* note 4, at 69 (“To be negligent, one does not advert to . . . the unreasonable risk that one is creating . . . It is adverting to such risks that converts one’s negligent conduct into recklessness.”).

rise to the core puzzle can be consistent because each claim concerns a different risk: negligent agents are aware of one risk and unaware of another.

What distinctive risk, then, do I propose is involved in negligence? Negligent conduct, as the Model Penal Code specifies, deviates from the standard of care a reasonable person would employ in perceiving risk.<sup>14</sup> That standard of care, I argue, governs inquiry: reasonable people gather evidence before acting because that evidence is required to evaluate the risks that possible actions would create. Failing to inquire, then, is risky. An individual who does not inquire before acting may fail to uncover evidence that her action would be unreasonably risky; by failing to inquire she risks inadvertently performing an unreasonably risky action and thereby causing harm. Negligent agents disregard that risk: they are epistemically reckless, or reckless with respect to how they inquire into evidence about their actions. They do not differ from reckless agents in why they are culpable for disregarding risk: each sort of agent possesses the same sort of consciousness or awareness of risk that is required for culpability. Rather, because reckless and negligent agents disregard risks that arise from different aspects of conduct, prohibitions on each guide individuals' conduct in different ways. Since recklessness involves the disregard of risks that arise from acting, prohibitions on recklessness prohibit acting; since negligence involves the disregard of risks that arise from failing to inquire, prohibitions on negligence mandate inquiry. And, I will argue, the requirement to inquire is different in important ways from the requirement not to act: because mandating inquiry interferes less with individual autonomy than prohibiting action does, the degree of risk required to justify prohibiting the failure to inquire is lower than the degree required to justify prohibiting action. Even though recklessness and negligence involve the same sort of culpable mental state, then, the law correctly distinguishes them as different kinds of mens rea when defining offenses.

Part I first explains how my approach to negligence is distinctive. While on the standard approach negligent agents are culpable for inadvertently creating the very same risk that they would be reckless for creating advertently, I instead take negligent agents to be culpable for failing to recognize that risk: negligence differs

---

<sup>14</sup> MODEL PENAL CODE § 2.02(2)(d) (AM. LAW INST. 1962).

from recklessness not in what psychological attitude grounds culpability but rather in what a defendant is culpable for. Part II analyzes when the failure to recognize a risk is negligent: because agents must sometimes gather evidence before acting in order to avoid inadvertently creating unreasonable risks, they may be culpable for failing to recognize a risk if they fail to inquire appropriately before acting. That failure to inquire, I argue, involves the disregard of a risk of a risk—namely, the risk of performing an action that would be judged unreasonably risky given the additional evidence that inquiry would have revealed. Part III then analyzes the culpability of criminal negligence, as I have defined it. Because I have identified a distinctive kind of risk involved in negligence, my account requires no innovations in the psychological basis of culpability: just as reckless agents consciously disregard risk, negligent agents are culpable for consciously disregarding the risk of failing to inquire. In some sense, this proposal interprets negligence as a form of recklessness. Nonetheless, I will argue that the law correctly distinguishes them as kinds of culpability because they should play a different role in defining offenses, given differences in the kinds of risk involved in each. Finally, I will argue that my approach, on which criminally negligent agents must possess the sort of subjective awareness that is generally required for criminal liability, better distinguishes criminal negligence from ordinary negligence than the standard approach.

### **I. Negligence as Culpable Ignorance of Risk**

Legal definitions of criminal negligence and of recklessness each require a culpable action to involve a similar kind of risk. The Model Penal Code, for example, defines both negligence and recklessness as involving “a substantial and unjustifiable risk that the material element [of an offense] exists or will result from [the defendant’s] conduct.”<sup>15</sup> And ordinary examples of reckless and negligent conduct create similar risks: to cut off another car on the highway, for example, risks harm in an intuitive sense both when the driver is recklessly aware of oncoming traffic and when he is negligently unaware of it.<sup>16</sup> Reckless and negligent conduct, then, do not

---

<sup>15</sup> MODEL PENAL CODE § 2.02(2)(c)–(d) (AM. LAW INST. 1962).

<sup>16</sup> To be sure, on some ways of understanding probability such actions might not be risky—if a particular instance of cutting off a car results in no harm, for example, then (assuming determinism) its probability of harm was zero given a full microphysical description of the state of affairs when it was performed. But there is nonetheless an

appear to differ with respect to the risk of harm created by the action the defendant performed. Instead, the law appears to distinguish the two grades of mens rea based on what sort of attitude the defendant held towards that risk: a reckless defendant “consciously disregards” that risk while a negligent defendant merely “should be aware of it.”<sup>17</sup>

These parallels and differences between the definitions of recklessness and negligence suggest one approach for understanding both what reckless and negligent conduct have in common and how their differences establish the lower culpability of the latter. Plausibly, individuals should avoid actions that create excessive risk. Inasmuch as negligent and reckless actions each create a substantial and unjustifiable risk of harm, then, that risk might constitute the reason why individuals should avoid acting either negligently or recklessly. Individuals should not cut other cars off on the highway, for example, because someone might get hurt; reckless and negligent defendants each fail to respond appropriately to that reason, since each kind of defendant performs an action of that type. The only difference between recklessness and negligence, then, would lie in the difference between the mental attitudes defendants of each sort take towards those risks: only reckless defendants, not negligent defendants, are aware of them. Theorists have often distinguished negligence from recklessness in this way. According to Kenneth Simons, “The academic debate over the propriety of criminal negligence liability ordinarily interprets negligence as unreasonable behavior in which the actor was unaware of the relevant risk, and interprets recklessness, knowledge, and intent as forms of culpability in which the actor was aware of the risk.”<sup>18</sup> Likewise, Doug Husak explains that “[T]he contested issue is easily described. Recklessness—which nearly everyone concedes to be a viable mode of culpability—is distinguished from

---

intuitive sense in which cutting off a car on the highway is risky—it is risky, one might say, under that description. For more on how best to construe the notion of probability in understanding offenses related to risk, see generally Eric A. Johnson, *Is the Idea of Objective Probability Incoherent?*, 29 LAW AND PHIL. 419 (2010).

<sup>17</sup> Compare MODEL PENAL CODE § 2.02(2)(c) with *id.* § 2.02(2)(d).

<sup>18</sup> Simons, *supra* note 10, at 371. Simons also ascribes this approach to the Model Penal Code. Kenneth W. Simons, *Dimensions of Negligence in Criminal and Tort Law*, 3 THEORETICAL INQUIRIES L. 283, 289–90 (2002) (“[R]ecklessness is defined (in part) as awareness that a harm may ensue or that an incriminating circumstance might obtain. At the same time, negligence is understood negatively, as a form of culpability in which the actor lacks such awareness. (Indeed, the only difference between negligence and recklessness under the Code is this difference in awareness.)”). To be clear, Simons’s own account of criminal negligence rejects what he takes to be the standard terms of the academic debate. See generally Simons, *supra* note 10.

negligence—which is highly controversial—by a single factor: the presence or absence of consciousness or awareness of a substantial and unjustifiable risk the defendant disregards.”<sup>19</sup> Negligent and reckless defendants are culpable for creating the same risk, on this understanding, but they are culpable in different ways for creating it.<sup>20</sup>

The core puzzle of criminal negligence follows straightforwardly from this characterization of the terms of debate. In paradigm examples of criminal conduct, run-of-the-mill instances of intentional wrongdoing, the wrongdoer’s culpability is established through some sort of awareness of the features of his action that he should have taken as reasons not to perform it. The intentional murderer, say, intends his action to cause the death of another. And his decision to act makes him culpable for the death that results precisely because, in acting, he was aware that he would cause that result: he recognized the reason in virtue of which acting would be impermissible, yet he acted nonetheless. The standard characterization of criminal negligence, though, makes it impossible to establish the culpability of negligent agents through this sort of strategy. For if the substantial and unjustifiable risk that an action creates constitutes the reason why a defendant should not perform it, then the culpability of negligent agents cannot depend on their recognition of that reason, since by definition they are unaware of that risk. The search for an alternative basis for culpability for negligence, then, is necessary only because of how negligence has typically been conceptualized. Ordinary, awareness-based culpability obviously cannot explain why an individual is culpable for performing an action with features that he fails to recognize. Thus, if the wrongfulness of negligence depends on the risk of harm created by the defendant’s conduct—the risk that he should have recognized but failed to recognize—some other basis must be identified to explain the culpability of negligence.

---

<sup>19</sup> Douglas Husak, *Negligence, Belief, Blame and Criminal Liability: The Special Case of Forgetting*, 5 CRIM. L. & PHIL. 199, 200 (2011). Husak, too, goes on to voice some skepticism about this account of the distinction between recklessness and negligence. *Id.* at 216 (“An even more radical possibility is that the means by which positive law distinguishes manslaughter from negligent homicide—which depends on whether defendants believe they have created a risk—is defective and should be replaced by some other device to mark the contrast.”).

<sup>20</sup> If a defendant’s conduct creates a risk of harm but he lacks any kind of psychological basis for culpability, then he might be held liable in tort alone, without being liable to criminal punishment. See Simons, *supra* note 18, at 286 (“[T]he ‘standard’ conception of negligence employed in tort law . . . consists in creating an unreasonable risk of physical harm to another . . .”).

Obviously, though, the analysis of criminal negligence might proceed along another avenue: since the need for an alternative basis of culpability for negligence depends on the analysis of criminal negligence as inadvertent risk-taking, one might instead defend criminal negligence not by expanding the basis of culpability but rather by reconsidering the analysis of negligence. The standard approach suggests that negligent agents fail to respond appropriately to a risk they do not recognize, then asks how individuals can be culpable for failing to be appropriately guided by a reason of which they are unaware. But an alternative characterization might exist instead of why negligent defendants fail to appropriately respond to the reasons that bear on their conduct.<sup>21</sup> And while such defendants' unawareness of the substantial and unjustifiable risk their conduct creates precludes an awareness-based account of why they are culpable for creating that risk, an alternative account of the reason negligent agents disregard in acting might be consistent with standard, awareness-based accounts of culpability, for negligent defendants might be aware of that alternative reason. No innovations in the theory of culpability would be required on this analysis to explain why negligence is culpable. I will attempt this strategy: in my view, criminal liability for negligence may be accommodated within existing theories of culpability by reinterpreting negligence. Negligence and recklessness should not be understood as involving different kinds of culpability for the creation of a substantial and unjustifiable risk. Instead, recklessness and negligence each involve the culpable disregard of a risk, but negligent agents culpably disregard a different kind of risk than reckless agents do.

The standard account's interpretation of negligence and recklessness as each involving the creation of the same sort of risk follows naturally from the first sentences of the Model Penal Code's definitions of each, which mention the same

---

<sup>21</sup> In a preemptive critique of approaches of this sort, Husak argues that because the dispute over criminal negligence concerns "whether and under what circumstances persons are to blame *for* their wrongs or breaches of duty. . . . [T]he inquiry is not advanced by positing additional wrongs that negligent defendants are alleged to commit." Husak, *supra* note 19, at 200. I agree, to be sure, that the puzzle of culpability for criminal negligence is not solved merely by identifying what is wrong with negligent conduct. But the standard analysis of criminal negligence employs a conception of the wrongdoing involved in negligence that by its very terms forecloses appeal to awareness-based accounts of criminal culpability. Thus, Husak overstates in claiming that the inquiry cannot even be advanced by positing an additional wrong committed by negligent defendants, for it may be easier to explain how a negligent defendant is culpable for a different wrong than for the wrong the standard analysis takes him to commit.



substantial and unjustifiable risk produced by the defendant's conduct and which differ only in what relationship to that risk they each ascribe to the defendant.<sup>22</sup> But those definitions each contain two sentences, the second of which distinguishes negligence from recklessness in a manner that on its face is quite different from the distinction drawn in the first sentences.<sup>23</sup> A reckless defendant disregards a risk that "must be of such a nature and degree that, considering the nature and purpose of the actor's conduct and the circumstances known to him, *its disregard* involves a gross deviation from the standard of *conduct* that a law-abiding person would observe in the actor's situation."<sup>24</sup> By contrast, a negligent defendant should have been aware of a risk that "must be of such a nature and degree that the actor's *failure to perceive it*, considering the nature and purpose of his conduct and the circumstances known to him, involves a gross deviation from the standard of *care* that a reasonable person would observe in the actor's situation."<sup>25</sup> While the first sentences each describe the same sort of risk, the second sentences describe the "nature and degree" of that risk in different ways: they distinguish recklessness from negligence by requiring reckless and negligent defendants to violate different standards that apply to different aspects of conduct.<sup>26</sup> These distinct standards governing recklessness and negligence suggest an alternative approach for distinguishing the two: since actions may violate different standards by possessing different kinds of features, the distinct standards that govern recklessness and negligence may indicate that an action must possess different kinds of features to be reckless or negligent.

---

<sup>22</sup> See MODEL PENAL CODE § 2.02(2)(c)–(d) (AM. LAW INST. 1962).

<sup>23</sup> The account of recklessness I defend in chapter 2 of this dissertation, of course, likewise assigns relatively more importance to the Model Penal Code's second sentence than to its first.

<sup>24</sup> MODEL PENAL CODE, § 2.02(2)(c) (emphasis added).

<sup>25</sup> *Id.* § 2.02(2)(d) (emphasis added).

<sup>26</sup> In discussing the Model Penal Code's characterizations of recklessness and negligence, Simons comments that "the *only* difference between negligence and recklessness under the Code is this difference in awareness" of risk. Simons, *supra* note 10, at 290. In my view, Simons is mistaken: the definition of negligence not only requires the agent to lack the sort of awareness present in recklessness but also requires that lack of awareness itself to constitute a violation of a standard of care governing the agent's perception of risk. No parallel requirement concerning awareness must be met for recklessness. Of course, Simons surely would not deny that this difference exists, were it pointed out, and his own account of negligence (and of culpable mental states more generally) self-consciously departs from the Model Penal Code's. But, at least in my view, his omission of this additional requirement from the Code's definition of negligence exemplifies a general neglect of this aspect of the definition among theorists of criminal law, despite what I see as its considerable potential for better understanding negligence.

How, then, do those standards differ? Reckless defendants grossly deviate from a “standard of conduct” in their “disregard” of the risk; negligent defendants grossly deviate from a “standard of care” in their “failure to perceive” that risk.<sup>27</sup> A standard of conduct, presumably, regulates what sorts of action individuals may perform, and a standard of conduct that governs the disregard of risk therefore identifies when performing a particular action would constitute a prohibited disregard of risk. Because actions violate this standard in virtue of the risks they create, awareness-based accounts of culpability would require defendants to be aware of those risks to be culpable for violating the standard. By contrast, the standard violated by negligence concerns care, not conduct, and it regulates the agent’s perception of risk, not his disregard of it. An agent violates this standard, then, not because his action created an excessive risk but rather because he failed to exercise adequate care in perceiving the risks his action might create. And if defendants must violate this standard to act negligently, then an account of culpability for negligence need only explain how defendants are culpable for failing to perceive a risk: in particular, since negligent defendants need not violate any standard regulating the disregard of risk, they are not culpable for disregarding the risks that they failed to perceive.

This interpretation might be implausible were the differences I identify between second sentences of the definitions of negligence and recklessness absent from or even contradicted by the first sentences. But in fact even the first sentences of each definition suggest, though more subtly, the distinction between negligence and recklessness that I propose. Certainly, both those sentences do define reckless and negligent conduct in terms of the relationship between the defendant and the substantial and unjustifiable risk that his conduct creates. But each specifies a different sort of relationship. Recklessness is defined purely in terms of a psychological state: the defendant must “consciously disregard” a risk.<sup>28</sup> Since recklessness is defined in terms of the defendant’s psychological attitude towards risk, that attitude can plausibly be understood as the source of his culpability for

---

<sup>27</sup> Compare MODEL PENAL CODE § 2.02(2)(c) with *id.* § 2.02(2)(d). In addition, the two standards differ in their stringency—conduct need only be law-abiding to avoid recklessness while care must be reasonable to avoid negligence, which plausibly sets a higher bar—though I will not focus on that distinction here.

<sup>28</sup> MODEL PENAL CODE § 2.02(2)(c).

creating the risk. But the definition of negligence is different. It characterizes the defendant's relationship to the risk normatively rather than psychologically: he acts negligently when he "should be aware" of that risk.<sup>29</sup> Since "conscious disregard" is merely descriptive, it can only be interpreted as explaining a defendant's culpability for violating some other standard—presumably, a rule prohibiting the creation of substantial and unjustifiable risks. But the definition of negligence itself describes the defendant as violating a rule, one that governs the perception of risk rather than its creation: he should have been aware of a risk when he was not, in fact, aware of it. Negligent defendants may be culpable for this violation. And if they are, that culpability itself justifies the imposition of criminal liability; there is no need to further explain how a defendant is also culpable for creating the risk he failed to perceive.

To be sure, because the definition of negligence, like that of recklessness, does refer to a substantial and unjustifiable risk, no account of negligence can disregard that risk entirely. Even if a defendant has failed to observe the standard of care of a reasonable person, that failure satisfies the definition of negligence only if it is further true that he should have been aware of a substantial and unjustifiable risk. A theory of negligence must explain this requirement. Accounts that take negligent defendants to be culpable for the risk created by their conduct may produce a simple and obvious explanation: the defendant cannot be culpable for inadvertently creating a risk unless his conduct actually does create it. But my account cannot explain the requirement in this way, since on my view negligent defendants are culpable not for the risk their conduct creates but rather for their failure to perceive it. Thus, I must advance a different explanation for the requirement: in my view, it relates to causation rather than culpability.

Ordinarily, individuals commit offenses defined in terms of negligence only when they negligently cause certain results, such as the death of another. Intuitively, such agents should be criminally liable only when their negligence is itself the cause of that result.<sup>30</sup> A negligent individual should not be punished if her conduct would

---

<sup>29</sup> MODEL PENAL CODE § 2.02(2)(d).

<sup>30</sup> The account of grading by result that I defend in chapter 1 of this dissertation explains why negligently causing an unintended result should be a crime only if the result would not have occurred had the defendant not been

have caused a harmful result even had she not been negligent—say, if a driver negligently fails to look before turning, even though in fact no car is coming, only to accidentally cause a freak accident when turning due to an unforeseeable mechanical failure unrelated to the failure to look. Thus, negligence is criminal only when a defendant’s negligence causes some change in the results of her conduct. And the failure to exercise adequate care can cause a change in those results only if a defendant would have acted differently but for her failure to exercise adequate care.<sup>31</sup> If someone who did exercise adequate care would not have perceived that acting would create a substantial and unjustifiable risk, then a negligent defendant would have had no reason to act differently had she not been negligent. Consequently, a violation of the standard of care causes a particular harmful result only if the defendant would have been aware of a substantial and unjustifiable risk but for her negligence. This constraint on which results the failure to exercise adequate care causes explains why negligence is defined to apply not simply because an agent violates the standard of care but rather only if, in addition, she should have been aware of a substantial and unjustifiable risk. But because this restriction arises from the requirement of causation, the substantial and unjustifiable risk created by a negligent defendant’s conduct need not explain her culpability; instead, she is culpable simply for violating the standard of care that governs the perception of risk.

This explanation for why defendants ordinarily may be negligent only when their conduct creates a substantial and unjustifiable risk may raise a different objection, however. On my view, the culpability of negligent defendants depends only on whether they comply with the standard of care observed by the reasonable person in perceiving risk. But individuals do not ordinarily commit a crime simply by violating that standard of care; instead, most offenses are defined also to require that violation to cause a particular harmful result. If my account of culpability for

---

negligent. On my view, grading by result guides conduct effectively because it leads offenders to expect a higher likelihood of punishment for offenses likelier to cause harm. Since an offender’s expectation of the likelihood of punishment guides him to avoid negligence, the law should calibrate that likelihood to match the actual strength of his reason to avoid it. But if a harmful result would occur regardless of whether he was negligent, then that result is not a reason not to be negligent, and the possibility of such results does not strengthen his reasons to avoid negligence. Consequently, in order to match the expected likelihood of punishment for negligence to the strength of the actual reasons to avoid negligence, such results cannot give rise to punishment, either.

<sup>31</sup> See, e.g., MODEL PENAL CODE § 2.03(1) (“Conduct is the cause of a result when: (a) it is an antecedent but for which the result in question would not have occurred . . .”).

negligence is correct, though, it may seem unclear why this restriction would exist: if all defendants who violate the standard of care governing the perception of risk are culpable, then why wouldn't the law define an offense that can be committed simply by violating that standard, rather than holding negligent defendants liable only if they also cause certain results? Why should a defendant's criminal liability depend on luck, rather than on his culpability alone?

If this objection were valid, however, it would apply well beyond negligence. Not only are offenses of negligence defined using results; individuals must cause certain results or engage in certain kinds of conduct to commit most crimes defined using any kind of mens rea.<sup>32</sup> Most offenses of recklessness, for example, are committed only if a defendant recklessly causes a particular result; because conviction for such crimes requires causation in addition to culpability, a defendant is not guilty simply for consciously disregarding a substantial and unjustifiable risk.<sup>33</sup> Thus, just as on my account actions may be culpably negligent without being criminal, so too may actions be culpably reckless without being criminal, if an agent consciously disregards a substantial and unjustifiable risk without satisfying the actus reus of any offense. To be sure, perhaps crimes should be defined so that a defendant's guilt cannot depend purely on his good or bad luck in the results of his culpable conduct. Philosophers have debated whether moral evaluations of conduct are sensitive to this sort of moral luck,<sup>34</sup> and many legal theorists argue against this method of defining crimes because they believe criminal punishment should not be sensitive to luck.<sup>35</sup> Those arguments, if correct, would disallow offenses of negligence, as I have understood them.<sup>36</sup> But by and large those arguments have not

---

<sup>32</sup> E.g., MODEL PENAL CODE § 210.1 ("A person is guilty of criminal homicide if he purposely, knowingly, recklessly or negligently causes the death of another human being.").

<sup>33</sup> E.g., MODEL PENAL CODE § 212.4(1) ("A person commits an offense if he knowingly or recklessly takes or entices any child under the age of 18 from the custody of its parent, guardian or other lawful custodian . . .").

<sup>34</sup> For classic treatments of the problem, see T. Nagel, *Moral Luck*, 50 PROC. ARISTOTELIAN SOC'Y 137 (1976); B.A.O. Williams, *Moral Luck*, 50 PROC. ARISTOTELIAN SOC'Y 115 (1976).

<sup>35</sup> E.g., Sanford H. Kadish, *Foreword: The Criminal Law and the Luck of the Draw*, 84 J. CRIM. L. & CRIMINOLOGY 679 (1994). Proposals for alternative methods of defining offenses tend more often to focus on the discrepancy between attempts and completed crimes than on the discrepancy between culpable risk-creation that does and does not cause harm. See, e.g., Joel Feinberg, *Equal Punishment for Failed Attempts: Some Bad but Instructive Arguments Against It*, 37 ARIZ. L. REV. 117, 119–21 (1995) (proposing to replace homicide offenses with a new offense of "wrongful homicidal behavior").

<sup>36</sup> I reject those arguments, at least with respect to unintended results, in chapter 1 of this dissertation.

been accepted: criminal liability does depend on harm in addition to culpability.<sup>37</sup> The role that causation plays in negligence, as I analyze it, is thus entirely consistent with the role that causation plays across criminal law.

Furthermore, were the arguments against the use of causation to determine liability to be accepted, my account of negligence could easily be modified to accommodate the change. Some offenses already omit a requirement to cause harm. The Model Penal Code’s general offense of reckless endangerment, for example, provides that “[a] person commits a misdemeanor if he recklessly engages in conduct which places or may place another person in danger of death or serious bodily injury;”<sup>38</sup> other offenses prohibit specific forms of reckless endangerment.<sup>39</sup> Analogous offenses of negligent endangerment, which would criminalize negligence even absent the causation of harm, are possible; some exist.<sup>40</sup> The law, then, could regulate risky activities solely through offenses that criminalize the mere creation of risk, regardless of whether it results in harm, though there is perhaps reason to doubt whether widespread prosecutions of risk-creation that fails to result in harm would be feasible, given that the occurrence of harm is often the only reason why the riskiness of the conduct that caused it is investigated at all. Under such a regime, offenses of negligence might all be replaced by offenses of negligent endangerment; alternatively, the requirement of endangerment could be omitted entirely, and culpable violations of the standard of care of the reasonable person might be criminalized on their own.

---

<sup>37</sup> *E.g.*, Kadish, *supra* note 35, at 670 (describing the “near universal acceptance in Western law” of the dependence of criminal liability on whether harm is caused); Nagel, *supra* note 34, at 140–43 (describing how commonsense moral evaluations of actions are sensitive to the lucky or unlucky results those actions cause).

<sup>38</sup> MODEL PENAL CODE § 211.2. Because a defendant commits this offense only when he “places or may place another person in danger” in addition to acting recklessly, even reckless endangerment may contain a causal element, but one requiring the causation of danger (however that is understood) rather than of harm.

<sup>39</sup> *E.g.*, MODEL PENAL CODE § 220.2(2) (“A person is guilty of a misdemeanor if he recklessly creates a risk of catastrophe in the employment of fire, explosives or other dangerous means listed in Subsection (1).”).

<sup>40</sup> *E.g.*, Clean Air Act of 1963 § 113(c)(4), 42 U.S.C. 7413(c)(4) (2018) (providing that “[a]ny person who negligently releases into the ambient air any hazardous air pollutant . . . and who at the time negligently places another person in imminent danger of death or serious bodily injury” commits an offense); MONT. CODE ANN. § 45-5-208(1) (2021) (“A person who negligently engages in conduct that creates a substantial risk of death or serious bodily injury to another commits the offense of negligent endangerment.”). Though these offenses do not require harm, they do require more than mere negligence—as with reckless endangerment, a defendant’s negligence must endanger another by risking death or serious bodily injury. Thus, these offenses do not criminalize culpability alone.

Thus, whether crimes of recklessness and negligence should require the causation of harm is largely independent of the adequacy of my account of culpability for negligence. While the standard account takes recklessness and negligence to be different forms of culpability for risk-creation, then analyzes the failure to perceive risk as a novel form of culpability, I propose instead that negligence be understood as culpability for the failure to perceive risk. Recklessness violates a standard governing the disregard of risk; negligence violates a standard governing the care taken in perceiving risks. But each of these standards are distinct from the elements that define the offenses that may be performed recklessly or negligently. On any account, recklessness and negligence each define a form of culpable risk-creation, which may then define offenses of either endangerment or harm-causing. Neither recklessness nor negligence analyze culpability directly in terms of the elements of the offenses they define, since each ties culpability to the risk a defendant creates rather than to those elements themselves. Thus, my proposal to understand negligent defendants as culpable for their failure to perceive risk cannot be rejected simply because that failure is distinct from the other elements of the offenses they perform, for reckless defendants too are culpable for a reason—their conscious disregard of a risk—that is similarly distinct from the elements of the offenses they perform.

## **II. Ignorance of Risk and the Failure to Inquire**

On my view, then, negligent defendants are culpable for their failure to perceive the risks their actions create. But the failure to perceive a risk obviously cannot always be a crime. The world is full of risks that nobody could possibly hope always to recognize, and thus, inevitably, individuals will sometimes fail to perceive a risk that their actions create. Such individuals should not be punished if those risks materialize in harm. An account that analyzes negligence as the failure to perceive a risk must explain when such failures can give rise to criminal liability, and when they cannot. In this respect, too, my analysis of negligence will differ from those that follow the standard approach in taking negligence to involve culpability for the disregard of an unperceived risk. Some such accounts, like mine, do focus on the

defendant's failure to perceive a risk that he should have perceived.<sup>41</sup> But they tend to frame their task as explaining how negligent defendants may be culpable for creating a risk that they fail to perceive.<sup>42</sup> Because they focus on identifying when the failure to perceive a risk can give rise to culpability for creating that risk, they analyze when a failure to recognize risk can justify attributing that risk to the agent in the same way that the choice to create a risk would justify a similar attribution.<sup>43</sup> And using that approach, they typically conclude that the culpability of a failure to perceive risk depends on which of the agent's motivational or evaluative attitudes was responsible for the failure.<sup>44</sup>

But there is another, more natural answer to the question of when a defendant should be culpable because of a failure to perceive risk—the criterion explicitly given by the Model Penal Code itself. A defendant may be convicted of negligence only if his “failure to perceive [the risk] . . . involves a gross deviation from the standard of care that a reasonable person would observe in the actor's situation.”<sup>45</sup> This requirement, clearly, presupposes that there is some standard of care that a reasonable person would follow in identifying the risks that his conduct might create. Furthermore, since the definition of negligence is satisfied only by gross deviations from that standard, it would seem that the law itself uses that standard to limit the

---

<sup>41</sup> STARK, *supra* note 10, at 226 (“[N]egligence is best understood in terms of the defendant's failure to form a belief that there was a specific risk attendant upon her behavior [ϕ].”); GARVEY, *supra* note 10, at 337 (analyzing negligence as “the culpable failure to exercise doxastic self-control, i.e., control over desires that influence the formation and awareness of one's beliefs”); *see also* TADROS, *supra* note 10, at 238 (“We can make some progress towards finding an answer to the subjective/objective debate [concerning recklessness] by considering different ways in which we might fail to form true beliefs.”).

<sup>42</sup> STARK, *supra* note 10, at 266 (“[N]egligence as failure of belief has three necessary components beyond the taking of a substantial, unjustified risk [i.e. wrongdoing] . . . .”); GARVEY, *supra* note 10, at 337 (“[W]hen, if at all, is retributive punishment for the inadvertent creation of a lethal risk warranted, and why?”).

<sup>43</sup> STARK, *supra* note 10, at 227 (“It is that absence of awareness of risk that links the defendant, as an agent, to her wrongdoing – her unjustified risktaking and its consequences.”); TADROS, *supra* note 10, at 244 (“[I]f we are responsible for those things that reflect on the agent qua agent, and the agent's character is constitutive of the agent qua agent, then we are responsible for our beliefs. For our beliefs, in manifesting the virtues and vices that we have, normally reflect on us qua agents.”); GARVEY, *supra* note 10, at 363 (“But if the culpability of an actor who unwittingly creates a lethal risk is not the culpability of choice, then what is it the culpability of? On the account offered here, the answer is failed self-control.”).

<sup>44</sup> STARK, *supra* note 10, at 250 (“[T]he feature of the defendant's character that prevents the formation of a belief that there is a risk attendant on ϕ-ing should give rise to a finding of culpability only if it has been accepted as being representative of, and in line with, the defendant's wider motivational framework.”); TADROS, *supra* note 10, at 264 (“In cases where the defendant acts in accordance with a mistaken belief, the criminal law is required to distinguish between agents whose beliefs manifest the kind of vice for which the attribution of criminal responsibility is appropriate and those that do not.”); GARVEY, *supra* note 10, at 381 (“Which desires the law should excuse an actor's failure to control will depend on the norms to which it defers, and when it is prepared to defer to them.”).

<sup>45</sup> MODEL PENAL CODE § 2.02(2)(d) (AM. LAW INST. 1962).



scope of criminal liability for negligence. If the failure to perceive a risk is negligent only when it grossly deviates from the standard of care of the reasonable person, then, that standard itself would be the most natural place to find a principle for distinguishing those failures to perceive risk that are criminal. What care, then, does the reasonable person take concerning the risks his conduct creates? By identifying that standard, we may identify when failures to perceive risk grossly deviate from it, and thus when they constitute criminal negligence. The principle limiting which failures to perceive risk are criminal need not turn on the psychological attitudes that suffice for culpability.

Much of belief-formation consists in evaluating and responding to evidence. Perhaps these inferential processes are not subject to our control, or at least not in the same way action is; perhaps we simply judge whether the evidence supports a particular belief, then believe, or not, upon that basis.<sup>46</sup> The care individuals take in identifying risk, then, may not directly influence what conclusions they infer from their evidence. But in addition to reasoning based on our evidence, we must also investigate the world to collect that evidence. Though some evidence collection is arguably automatic—perception, for example—what evidence we possess also depends to a substantial extent on our voluntary decisions to inquire into the world around us. In Gideon Rosen’s description:

The active side of epistemic life is inquiry, a process in which we ask questions, formulate hypotheses, gather evidence, perform experiments, reflect on our results, double-check our calculations, and so on. As we take these active measures (some of which involve overt behavior, others of which are purely mental), we revise our views in light of our findings, and this process of belief revision is for the most part a passive matter. After all the active steps have been taken, one simply finds oneself believing what one believes . . . .<sup>47</sup>

---

<sup>46</sup> There is, to be sure, a large literature on the question, often described as one of the “ethics of belief,” of whether moral considerations may sometimes oblige or permit our beliefs to deviate from what the evidence alone would support. *See generally* Andrew Chignell, *The Ethics of Belief*, STAN. ENCYCLOPEDIA PHIL. (Edward N. Zalta ed., Spring 2018 ed.), <https://plato.stanford.edu/archives/spr2018/entries/ethics-belief/>. I do not mean my comments here to endorse a position in those debates, only to acknowledge that it is a difficult question whether moral considerations may properly influence whether evidence supports a belief.

<sup>47</sup> Rosen, *supra* note 12, at 600.

Even if in managing our beliefs we must simply take as given our dispositions to infer certain conclusions from our evidence, the choices we make about how to investigate the world—choices that determine what evidence is available to employ in inference—are simply choices to perform a particular kind of action. Just as we must comply with rules when acting in general, so too can certain rules govern inquiry, and just as an agent can be culpable for failing to act as he ought to, so too can he be culpable for failing to inquire as he ought to.

Obvious reasons exist why an agent ought sometimes to inquire before acting in order to improve the quality of her evidence. Probability estimates generally depend on the information that is employed in making them. Thus, an individual's estimates of the likelihood that a particular action will cause a particular harm will depend on what information she possesses about the action and its circumstances. Furthermore, as she learns additional information about that action, her estimates of its risk of causing harm will typically become more accurate. To estimate probabilities, we generally identify the fraction of possibilities in which a particular proposition is true: to say that an action has a particular probability of causing harm is to judge that harm occurs in that fraction of all the possible states of affairs in which it is performed.<sup>48</sup> Such judgments are relevant to action because we do not know which possible state of affairs is actual: if we have no basis for thinking that any such possibility is likelier than any other, then our expectation of whether acting will cause harm must be the fraction of all possibilities in which harm occurs.<sup>49</sup> Which states of affairs are possible, in turn, depends on the evidence we possess—in particular, states of affairs are possible if they are consistent with our evidence. As we gain additional evidence, then, we can refine our probability estimates by eliminating from consideration any possibilities incompatible with our new evidence. And eliminating such impossible scenarios from our reasoning will typically improve the decisions we make. Those scenarios were impossible all along, since the evidence excluding them

---

<sup>48</sup> *E.g.*, FRANK H. KNIGHT, *RISK, UNCERTAINTY, AND PROFIT* 220 (1921) (“If the chance of any particular result is more or less than one half, it is held to be axiomatic that there is a greater number of possible alternatives which yield this result [or do not yield it] than of the other kind . . .”).

<sup>49</sup> *Id.* (“The whole mathematical theory of probability is obviously a simple application of the principles of permutations and combinations for finding out the number of alternatives. Absolute indifference between the alternatives is taken for granted.”).

existed even before we learned it. And the results acting would produce in scenarios that are not actual should not influence our behavior in the actual world. We cannot exclude all impossible scenarios from our reasoning, since perfect information is beyond our reach. But when we can exclude some such scenarios, we can make better decisions on the basis of the probability judgments that result.

Risky activities constitute one context in which improving our probability estimates may affect how we ought to act.<sup>50</sup> In general, individuals ought not perform actions that impose excessive risks on others. But any individual deliberating over a possible action can only employ the evidence she actually possesses in evaluating it. And since an evaluation of an action may depend on what information is employed in estimating its risk, whether it must be avoided because of that risk will depend on that information. In particular, in some cases an action may create a tolerable degree of risk if evaluated using the evidence that an individual actually possesses, even though its risk would be excessive were it evaluated using additional evidence that she might obtain. Such actions ideally would not be performed. Since the scenarios incompatible with that additional evidence cannot obtain, they are not relevant to whether the benefits of that action are outweighed by its possible costs; instead, it could be more accurately evaluated by considering only the possible costs and benefits in scenarios compatible with that additional evidence. But, obviously, in order to consider that evidence when evaluating the action any agent must actually engage in inquiry and uncover it before deciding what to do. Before a driver turns to look for oncoming traffic, for example, the risk of turning might be low if evaluated using the information she already possesses concerning the time of day, the activity at the intersection, and so forth. But by looking, she will gain key additional information—namely, whether any cars are coming. That evidence may dramatically affect the risk of turning. And the estimate made without that additional evidence is obviously less relevant to whether turning would be permissible: if a car is coming, it does not matter that turning would be safe in possibilities where the road is clear.

---

<sup>50</sup> Reducing the risks that actions will harm others is not the sole consideration governing inquiry; other reasons may sometimes supplement (and perhaps sometimes contradict) it. Perhaps knowledge is an intrinsic good that we have moral reason to pursue it for its own sake. My account of criminal negligence, though, will treat the standard of care that governs criminal negligence as concerned only with avoiding the risks of acting in ignorance.

Failing to inquire, then, is itself a form of risky behavior. In some cases, were individuals to inquire before acting they would identify additional evidence that substantially increases the likelihood that an action will cause harm; armed with that extra evidence, they could either avoid acting or modify their conduct to reduce the risk. Of course, in many cases that additional evidence will not exist, and when it does not exist omitting inquiry will cause no harm. But by omitting inquiry individuals take the risk that their omission will prevent them from acquiring information that would lead them to modify their behavior. In some sense, this risk is perhaps better understood as a risk of a risk: even when the relevant evidence does exist, harm still might not occur, since that evidence will ordinarily indicate only that acting would be risky, not that it would necessarily be harmful. But a risk of a risk is still a risk: since riskier actions more often do cause harm, individuals who take the risk that their actions are riskier than they realize will cause harm more frequently, in the aggregate, than individuals who reduce that risk by inquiring before they act. And since the likelihood of harm when inquiry is omitted is higher than when it is performed, the omission itself increases the risk of harm. By omitting inquiry, individuals risk depriving themselves of evidence they would use to avoid inadvertently performing an action with a substantial and unjustified risk of harm. By failing to look before turning, for example, a driver risks causing an accident because he did not realize there was an oncoming car.

Reasonable people are sensitive to the risk of harm to others that their actions create, and they take reasonable steps to avoid such risks in their behavior. Omitting inquiry is simply another way to create risk. Thus, the reasonable person avoids those risks by taking reasonable steps to inquire before acting.<sup>51</sup> This principle, in my view, determines when the failure to perceive the risk created by one's conduct constitutes criminal negligence: that failure must have involved a gross deviation from the standard the reasonable person would observe in inquiring into the risks that an action might create. When an agent's omission of inquiry grossly deviates from that

---

<sup>51</sup> Other theorists have made similar points: for example, Alexander Sarch argues that the primary duty to avoid criminal conduct gives rise to a secondary duty to engage in reasonable investigation, though because his focus is willful ignorance he applies that doctrine in cases when a defendant possesses "substantial confidence" that a contemplated action will constitute a crime. ALEXANDER SARCH, CRIMINALLY IGNORANT 114–15 (2019).

standard, and he would have identified evidence that his action would create a substantial and unjustifiable risk had he inquired in accordance with that standard, then in a natural sense he should have been aware of the risk: he would have been aware of it had he taken adequate care in inquiring before acting. On my account, such agents are criminally negligent. Whether the failure to perceive a risk is a crime depends on whether it results from omitting inquiry that the reasonable person would have performed.

Obviously, the law cannot require individuals to engage in all forms of inquiry that are available; the risks created by omitting inquiry, like all risks, can be limited but cannot be eliminated completely. Inquiry must have limits: we cannot acquire all the evidence that exists but rather must at some point act on the limited information that we possess. (Indeed, since delay carries risks of its own at some point further inquiry would itself be unreasonable.) Thus, not all omissions of inquiry grossly deviate from the standard of the reasonable person: the key question is whether the defendant has inquired adequately, not whether he has inquired at all. (When I speak of the failure to inquire in what follows, I will always mean adequate inquiry.) As with all risks, whether the risk created by omitting further inquiry may be ignored depends on the magnitude of that risk and the cost of avoiding it. If inquiry would likely reveal additional evidence that would dramatically increase the probability of an action's causing harm, omitting inquiry would be very risky; if not, omitting it would create only a small risk. And the more difficult inquiry would be, the higher the cost of avoiding the risk.

Articulating these considerations in the abstract leaves open how they apply in particular cases; as Rosen puts it, “the epistemic requirements that govern inquiry . . . vary massively from case to case, and it is hard to say anything very general about their content.”<sup>52</sup> But the law itself is no more specific: the Model Penal Code defines negligence simply as a failure to perceive risk that grossly deviates from the standard of care of the reasonable person. Individuals bound by the law and juries applying it in adjudication must evaluate whether any particular omission grossly deviated from that standard based on the facts of that particular case. Some omissions may be easy

---

<sup>52</sup> Rosen, *supra* note 12, at 601.

to evaluate: drivers should look before turning, but they need not conduct a full mechanical inspection of their vehicles before every occasion on which they drive. Others may be harder to evaluate. But many doctrines in criminal law, from recklessness to duress, define prohibited conduct through irreducibly normative standards that are highly sensitive to the particular facts of individual cases.<sup>53</sup> Criminal negligence depends on a similar standard, in my view—one that governs when individuals must inquire before acting.

On my account, negligence is in some respects similar to a different kind of culpability—willful ignorance. A defendant who fails to know some fact but who is willfully ignorant of it may be convicted of a crime defined to require knowledge of that fact.<sup>54</sup> Agreement does not exist concerning why willfully ignorant defendants are as culpable as knowing ones.<sup>55</sup> But some proposals explain the culpability of willful ignorance through the defendant’s failure to inquire. Alexander Sarch, for example, argues that willfully ignorant defendants’ breach of a duty to reasonably inform oneself makes them as culpable as knowing defendants.<sup>56</sup> Similarly, Gideon Yaffe argues that the failure to inquire reveals willfully ignorant defendants to be just as willing as knowing defendants to disregard the interests of others in their conduct, and thus to be equally culpable.<sup>57</sup> Since negligence, obviously, is less culpable than knowing wrongdoing, an account of negligence must identify a kind of culpability that is less serious than knowing wrongdoing. If individuals who fail to inquire were as culpable as knowing wrongdoers, though, the failure to inquire could not constitute criminal negligence. Thus, I must explain why the failures to inquire that constitute negligence are less culpable than those that constitute willful ignorance.

---

<sup>53</sup> *E.g.*, MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962) (defining recklessness to require “a gross deviation from the standard of conduct that a law-abiding person would observe in the actor’s situation”); *id.* § 2.09(1) (defining duress to require “the use of, or a threat to use, unlawful force against [the defendant’s] person or the person of another, that a person of reasonable firmness in his situation would have been unable to resist).

<sup>54</sup> *See, e.g.*, United States v. Heredia, 483 F.3d 913, 918 (9th Cir. 2007) (en banc) (“‘[K]nowingly’ in criminal statutes . . . includes the state of mind of one who does not possess positive knowledge only because he consciously avoided it.” (first alteration in original) (quoting United States v. Jewell, 532 F.2d 697, 702 (9th Cir. 1976) (en banc))).

<sup>55</sup> Douglas N. Husak & Craig A. Callender, *Willful Ignorance, Knowledge, and the “Equal Culpability” Thesis: A Study of the Deeper Significance of the Principle of Legality*, 1994 WIS. L. REV. 29–70, 53–58 (1994) (discussing the “equal culpability” thesis).

<sup>56</sup> SARCH, *supra* note 51, at 110.

<sup>57</sup> Gideon Yaffe, *The Point of Mens Rea: The Case of Willful Ignorance*, 12 CRIM. L. & PHIL. 19, 29–31 (2018).

A mere failure to inquire does not constitute willful ignorance on its own. Instead, a willfully ignorant defendant may be convicted of an offense requiring knowledge only if he fails to inquire and also satisfies two additional conditions. Because each of those conditions plausibly enhances his culpability, they explain in part why willfully ignorant defendants are as culpable as knowing ones. Even though willful ignorance is considerably more culpable than negligence, then, that culpability does not arise solely from the failure to inquire: willfully ignorant defendants are culpable to a higher degree both because they fail to inquire and because they satisfy the other conditions required for willful ignorance. Consequently, the example of willful ignorance does not show that the failure to inquire, on its own, gives rise to a degree of culpability incompatible with my analysis of criminal negligence in terms of the failure to inquire. Instead, because the greater culpability of willfully ignorant defendants depends on their additional satisfaction of these other conditions, willfully ignorant defendants may be as culpable as knowing ones while defendants who merely fail to inquire, without also satisfying those conditions, may be much less culpable. Thus, so long as criminally negligent defendants do not satisfy those two conditions, the equal culpability of willfully ignorant defendants and knowing defendants is perfectly compatible with my account of criminal negligence.

First, a willfully ignorant defendant must be aware of a high risk that his action will constitute a crime. The Supreme Court, for example, explains that “[t]he defendant must subjectively believe that there is a high probability that a fact exists” for him to be willfully ignorant of that fact;<sup>58</sup> similarly, the Model Penal Code’s provision governing willful ignorance provides that “[w]hen knowledge of the existence of a particular fact is an element of an offense, such knowledge is established if a person is aware of a high probability of its existence.”<sup>59</sup> Since the awareness of a substantial and unjustifiable risk ordinarily suffices for recklessness, it is natural to understand this condition as requiring willfully ignorant defendants to be reckless independent of their failure to inquire. Sarch, for example, explicitly requires the willfully ignorant defendant to be reckless independently of his failure to

---

<sup>58</sup> *Global-Tech Appliances, Inc. v. SEB S.A.*, 563 U.S. 754, 769 (2011). Although *Global-Tech* was a patents case, not a criminal case, the Court imported the doctrine of willful blindness from criminal law to resolve it. *Id.*

<sup>59</sup> MODEL PENAL CODE § 2.02(7); *see also id.* cmt. 9.

inquire.<sup>60</sup> But, obviously, if a willfully ignorant defendant must recklessly disregard a subjective belief that his action is highly likely to constitute an offense, that the level of culpability of that recklessness on its own exceeds the culpability level of criminal negligence. The failure to inquire may further increase his culpability to be equivalent to negligence: indeed, Sarch advances this exact analysis of willful ignorance.<sup>61</sup> But the fact that a reckless defendant is more culpable than a negligent one if he also fails to inquire does not show that a defendant who fails to inquire without otherwise being reckless must have a level of culpability incompatible with criminal negligence.

Second, willful ignorance—as the name suggests—requires willfulness. The Ninth Circuit pithily explains that willfully ignorant defendants “don’t know because they don’t want to know”;<sup>62</sup> more precisely, the Supreme Court explains that, “the defendant must take deliberate actions to avoid learning of that fact” of which he was ignorant.<sup>63</sup> Thus, both Sarch and Yaffe understand willful ignorance to apply only if the desire to avoid knowledge plays a particular role in the agent’s motivation for failing to inquire.<sup>64</sup> But, obviously, ignorance is only one possible motivation for an agent’s failure to inquire. An individual’s inadequate concern for whether his action will inadvertently threaten others may lead him not to inquire even when he is not actively seeking to remain ignorant—perhaps he does not wish to bear the costs of

---

<sup>60</sup> SARCH, *supra* note 51, at 129 (describing the “two components of willful ignorance” as “breaching [the duty to reasonably inform oneself] and recklessly performing the actus reus while suspecting [the associated inculpatory proposition]”).

<sup>61</sup> *Id.* at 117–29 (defending and applying the “Added Culpability Thesis”). Unlike Sarch’s, Yaffe’s argument for why willful ignorance is as culpable as knowledge does not require the willfully ignorant defendant to recklessly disregard a risk that his action will cause harm; instead, the failure to inquire on its own establishes that a willfully ignorant agent’s disregard for the welfare of others is equivalent to that of the knowing wrongdoer. Yaffe, *supra* note 57, at 29–31. Thus, it might seem that on Yaffe’s account those who fail to inquire should be just as culpable as knowing wrongdoers regardless of whether they are independently reckless. Nonetheless, since Yaffe acknowledges that his argument does not apply when the agent assigns a low probability to the illegality of the act, even on his account the failure to inquire is not always as culpable as knowing wrongdoing. Furthermore, since his account does rely on the second additional doctrinal requirement for willful ignorance, *see infra* note 64, it does not show that the failure to inquire is as culpable as knowing wrongdoing.

<sup>62</sup> United States v. Heredia, 483 F.3d 913, 918 (9th Cir. 2007).

<sup>63</sup> Global-Tech Appliances, Inc. v. SEB S.A., 563 U.S. 754, 769 (2011).

<sup>64</sup> Sarch interprets this requirement as restricting willful ignorance to defendants whose purpose in failing to inquire is to remain ignorant, though he argues that the vague language some courts use to characterize willful ignorance may extend it to defendants who merely know that they will remain ignorant by failing to inquire, as well. SARCH, *supra* note 51, at 20–21. Yaffe does not directly understand the defendant’s attitude towards his failure to inquire in terms of standard culpability grades of the Model Penal Code; instead, he holds defendants to be willfully ignorant only if they fail to inquire because they would be unwilling to act were they to discover that their action constituted a crime. Yaffe, *supra* note 57, at 23. Willfully ignorant defendants, then, must be instrumentally motivated to remain ignorant, since were they not to remain ignorant they would be unwilling to act.



inquiring further, does not wish to delay acting, or simply cannot be bothered to care. And an agent's motivations for performing a wrongful action ordinarily affect his degree of criminal culpability: someone motivated to perform a wrongful action by the features that make it wrongful is more culpable than someone who performs it because he simply does not care enough about avoiding actions with those features.<sup>65</sup> The Model Penal Code identifies purpose as the most culpable grade of mens rea: someone who knows his action will cause a particular result is less culpable than someone whose purpose in acting is to cause that result.<sup>66</sup> And Yaffe and Sarch each appeal to the willfully ignorant defendant's motivations for remaining ignorant to explain his equal culpability to the knowing one.<sup>67</sup> But if the willfully ignorant defendant's motivation to remain ignorant explains his high degree of culpability for his failure to inquire, then agents who fail to inquire for other reasons need not be so culpable. Thus, the equal culpability of willfully ignorant and knowing defendants

---

<sup>65</sup> For the view that an important moral difference exists between aiming at a harmful result and merely failing to treat it as a sufficiently strong reason against acting, see Thomas Nagel, *The Limits of Objectivity*, in THE TANNER LECTURES ON HUMAN VALUES 77, 129–35 (1980). For applications of this distinction to criminal law, see R.A. Duff, *Criminalizing Endangerment*, in DEFINING CRIMES: ESSAYS ON THE SPECIAL PART OF THE CRIMINAL LAW 43 (R.A. Duff & Stuart Green eds., 2005); and Jeremy Horder, *The Classification of Crimes and the Special Part of the Criminal Law*, in DEFINING CRIMES: ESSAYS ON THE SPECIAL PART OF THE CRIMINAL LAW 21 (R.A. Duff & Stuart Green eds., 2005).

<sup>66</sup> E.g., MODEL PENAL CODE 2.02(5) (AM. LAW INST. 1962) (providing that the mental state of purpose suffices to commit an offense defined with the mens rea of knowledge, but not vice versa).

<sup>67</sup> Yaffe's argument infers that the willfully ignorant defendant assigns the same value as a knowing defendant to the interests of others from the fact that a willfully ignorant defendant's motivation for remaining ignorant is to avoid obstacles that might prevent him from performing a potentially illegal action. Yaffe, *supra* note 57, at 29–31. Obviously, this argument does not apply when a defendant fails to inquire for some other reason, such as the costs or delay of inquiry. Indeed, Yaffe explicitly notes that his argument assumes inquiry to be costless; thus, his argument simply does not apply to those who avoid inquiry because of its costs (broadly construed). *Id.* at 24. Sarch agrees that a defendant's mens rea with respect to his failure to inquire affects his degree of culpability. SARCH, *supra* note 51, at 122. Consequently, he accepts that those whose failure to inquire is merely reckless are less culpable than willfully ignorant defendants, ceteris paribus. *Id.* at 184–85. Nonetheless, he argues that reckless indifference can sometimes elevate a reckless defendant's culpability to the level of a knowing wrongdoer: in particular, he argues that iterated reckless ignorance may be as culpable as knowing ignorance. *Id.* at 186. One might accept that the failure to inquire is sometimes a source of culpability without accepting that iterating that failure increases culpability in the simple, aggregative manner that Sarch suggests. Furthermore, as Sarch acknowledges, courts do not generally accept this extension of the willful ignorance doctrine. *Id.* at 20 & n.65. But even if his argument is correct that this form of reckless ignorance is as culpable as knowing wrongdoing, that argument requires defendants to satisfy two additional conditions that ordinarily will not be satisfied by defendants that my theory considers criminally negligent. First, iterated reckless ignorance must obviously be iterated. *Id.* at 185. Second, Sarch takes iterated reckless ignorance to elevate a defendant's culpability to the level of knowledge only if he was already reckless: the defendant must have consciously suspected that his action would satisfy the elements of an offense. *Id.* at 186. Thus, even if reckless ignorance is as culpable as knowing wrongdoing given these conditions, the failure to inquire may be culpable only to the degree of negligence when those conditions are not met.

does not preclude defendants whose failure to inquire is not willful from having a degree of culpability that is consistent with criminal negligence.<sup>68</sup>

### **III. Criminal Negligence as Epistemic Recklessness**

Negligence challenges standard theories of culpability, I have argued, because the mental state defendants must lack to be negligent—awareness of the risk their actions create—is the very same mental state they must possess to be culpable. Culpability for negligence, then, would require a defendant to possess incompatible attitudes—both advertence and in advertence towards a single risk. But this puzzle rests on a particular conception of what negligence is: the standard account presupposes that negligent agents are culpable for creating the same risk that reckless agents create, and thus it analyzes negligence as a novel form of culpability for that very risk that the negligent agent fails to perceive. By contrast, I have argued that negligence involves an additional kind of risk—not the risk that by acting an individual will unintentionally cause harm but rather the risk that by failing to inquire an individual will fail to uncover evidence he might have used to better estimate, and thus avoid, the risk his action created.

Distinguishing these two risks allows my account to resolve the core puzzle of negligence without introducing into criminal law a novel form of culpability grounded in a mental state different from those involved in other grades of mens rea. I will argue that criminal negligence is epistemic recklessness, a form of culpability grounded in the same kind of mental state as recklessness but nonetheless distinct from recklessness in its character and in its degree of culpability. Having presented my account, I will then reply to the objection that, whatever its merits as a form of culpability, it fails to describe the notion of negligence actually employed in criminal law: in my view, it does adequately characterize when negligence is liable to criminal punishment, and it explains why only some forms of negligence are criminally

---

<sup>68</sup> If willful ignorance increases a defendant's culpability from what it otherwise would be, such that a willfully ignorant reckless defendant is as culpable as a knowing one, one might wonder whether willful ignorance can similarly increase the culpability of a merely negligent defendant. Intuitively, it seems to me that it can: someone who fails to inquire in order to avoid learning evidence that his action would be reckless seems to be to be just as culpable as a reckless defendant. And I think that Yaffe's account of willful ignorance could be more or less directly extended to justify this conclusion, though I will not extend it here. Nonetheless, this question lies outside the scope of my arguments in this chapter, and I will not further defend my intuitions about such cases.

punishable better than rival views do. Finally, I will compare my account with the tracing approach to negligence, which is in some ways similar but which differs in certain important respects. Because of those differences, my account can avoid the prominent objections levied against tracing.

*A. The Culpability of Criminal Negligence*

If culpability requires awareness, negligent agents must both advert and fail to advert to risk. If negligence involved only a single risk, these attitudes would be inconsistent. On my view, however, negligence involves two distinct risks. A negligent agent may both advert and fail to advert to risk, then, because his advertence and inadvertence concern different risks. Because he does not engage in inquiry, he does not perceive the substantial and unjustifiable risk that his action creates—that is, he fails to judge it risky to the degree that he would have had he acquired additional evidence through inquiry. He therefore creates that risk inadvertently: because he is unaware of it, it does not explain why he is culpable. One of the two attitudes constitutive of negligence—the agent’s inadvertence towards a risk—is directed towards this risk. But, in addition, another risk is involved in the negligent agent’s conduct—the risk that by failing to inquire he will fail to identify evidence that he could have used to more effectively reduce the risks his actions create. In some cases, the existence of such risks provides individuals with sufficient reason why they must inquire before acting. And while a negligent defendant’s failure to engage in inquiry prevents him from recognizing the greater risk he would have recognized, given the additional evidence that inquiry would have uncovered, no similar obstacle prevents him from being aware of the risks attendant on the failure to inquire itself. Thus, though the negligent defendant by definition cannot advert to the risk he would have perceived were he not negligent, he can advert to the distinct risk of failing to inquire before acting. This risk can explain his culpability: he culpably fails to inquire before acting. And because he does advert to that risk, no novel account of culpability is required to explain his culpability: disregarding the risk that by failing to inquire he will unintentionally perform an action more dangerous than he realizes may subject him to the same sort of culpability as the disregard of any other sort of risk.

The culpability of negligence and of recklessness, then, rests on the same psychological basis: negligent and reckless agents are each aware of and disregard the reason why they should not have acted as they did. Furthermore, in each case that reason involves a risk that their actions will cause harm. The negligent agent, like the reckless one, is culpable because he consciously disregards a risk. The distinction between reckless and negligent agents does not lie in some distinction between the type of mental state that undergirds their culpability; instead, like defendants guilty of intentional offenses, both reckless and negligent defendants are culpable because they are aware of considerations that should but do not govern their choice of how to act. Just as knowing and reckless defendants are distinct because they disregard different normative considerations of which they are each aware—knowing defendants are aware of and disregard the practical certainty that an action will have a particular feature, whereas reckless defendants are aware of and disregard the risk that it will have that feature—reckless and negligent defendants are distinct because they disregard different normative considerations of which they are each aware—namely, different risks. The reckless defendant disregards the unacceptably high risk that performing an action will cause harm; the negligent defendant disregards the unacceptably high risk that failing to inquire will cause harm by depriving him of evidence through which he could avoid performing an unreasonably risky action. Because my account identifies a risk that the negligent agent creates that is distinct from the risk he fails to perceive, it need not identify how culpability could depend on a different sort of mental state in order to explain how he may be culpable for creating a risk he fails to perceive.

On my account, then, negligence is in some sense a kind of recklessness, since negligent defendants are culpable for consciously disregarding a risk.<sup>69</sup> Negligence is epistemic recklessness: the culpable disregard of the risk of being inadequately informed about the risks of an action. Thus, violating the standard of care of the reasonable person is not the only criterion a defendant must satisfy to be negligent. Certainly, he must have failed to inquire in circumstances where the reasonable

---

<sup>69</sup> Thus, the mental state I ascribe to negligent defendants differs from the mental state of willfully ignorant defendants, who are willfully rather than recklessly ignorant. See *supra* notes 62–68 and accompanying text.

person would have inquired. But criminally negligent defendants must additionally be culpable for that failure to inquire, which in turn depends on whether they consciously disregard the risk created by their deviation from the standard of care of the reasonable person. Thus, even if from an objective standpoint an agent ought to have inquired before acting—if, say, evidence in fact existed that he could have identified and that would have led him to avoid creating a risk—he will not always be culpable for that failure. In particular, he would not be negligent if he failed to recognize and thus to consciously disregard the risk created by that failure, just as an individual who performs an objectively risky action may not be reckless if he, too, failed to recognize and thus to consciously disregard the risk created by his action.

Indeed, it perhaps should not be surprising for the same subjective element to determine whether an agent is culpable for violating the standards imposed on individuals by recklessness and negligence. The Model Penal Code's provisions concerning recklessness and negligence differ in how they define those standards, speaking of a standard of conduct or of care that each applies to different aspects of an individual's behavior.<sup>70</sup> But those provisions also contain a subjective element: that is, they each specify which of the agent's mental states must be considered in determining whether he has violated the standard. And—crucially—those two elements do not differ at all: the definitions of recklessness and negligence each evaluate the defendant “considering the nature and purpose of his conduct and the circumstances known to him.”<sup>71</sup> Thus, whatever differences may exist between reckless and negligent defendants, the Model Penal Code suggests that those differences do not lie in the kind of mental states or psychological attitudes responsible determining whether each sort of defendant is culpable.<sup>72</sup> My approach follows the Model Penal Code in this respect. In my view, recklessness and negligence impose standards that regulate different aspects of a defendant's behavior,

---

<sup>70</sup> Compare MODEL PENAL CODE § 2.02(2)(c) (AM. LAW INST. 1962) with *id.* § 2.02(2)(d).

<sup>71</sup> *Id.* § 2.02(2)(d). The corresponding provision in the definition of recklessness differs only in replacing “his” with “the actor's.” *Id.* § 2.02(2)(c).

<sup>72</sup> While an agent's purpose and his knowledge of circumstances are clearly kinds of mental states, the “nature . . . of his conduct” might be not obviously mental. Nonetheless, the Model Penal Code defines “conduct” as “an action or omission and its accompanying state of mind.” *Id.* § 1.13(5). Thus, the nature of conduct encompasses the mental state accompanying an action or omission.

but they do not differ in the sort of mental state that determines whether a defendant has violated the relevant standard.

Because this account takes culpability for negligence to depend on the same mental state that secures culpability for recklessness, it can be applied to particular cases only given an account of what is required for an individual to consciously disregard a risk. If negligence is in some sense a species of recklessness, then a full account of negligence requires an account of recklessness, too. But because I do not mean to tie my account of negligence to any particular account of recklessness, I will not here seek to combine it with any particular account of when an individual recklessly disregards a risk.<sup>73</sup> Instead, my conception of negligence is broadly compatible with multiple different accounts of recklessness. In this chapter, I have attempted to reconceptualize the notion of criminal negligence—to argue that it may be incorporated as a grade of culpability in criminal law without expanding the set of mental states that the law would otherwise treat as a basis for criminal liability. The kind of mental state that suffices for higher grades of culpability suffices for negligence, too. What instead distinguishes negligence from recklessness is the object of that mental state—negligence involves a different kind of risk, not a different kind of mental state directed at the same risk.

If my analysis of the structure of negligence may be combined with any account of the mental state required for recklessness, defendants will be guilty of negligence whenever they possess that attitude—whatever it may be—concerning the risk created by failing to inquire. If an explicit belief directed towards a risk that results in harm is required for recklessness, then to be negligent defendants must possess a belief of that sort about the risk created by failing to inquire.<sup>74</sup> If, instead, a wider range of psychological attitudes may render an individual culpable for creating a risk, then those attitudes may render a defendant culpable for the risks created by

---

<sup>73</sup> In chapter 2 of this dissertation, of course, I have developed an account of recklessness.

<sup>74</sup> *E.g.*, ALEXANDER, FERZAN, & MORSE, *supra* note 4, at 28 (defining recklessness in terms of the defendant's beliefs about the level of risk he is imposing); STARK, *supra* note 10, at 93 (defining recklessness in terms of the defendant's belief that a specific risk exists); Kimberly Kessler Ferzan, *Opaque Recklessness*, 91 J. CRIM. L. & CRIMINOLOGY 597, 638 (2001) (arguing that individuals may be responsible only if they are "introspectively aware" in a manner that "trigger[s their] speech center"); Jenny McEwan & St. John Robilliard, *Recklessness: The House of Lords and the Criminal Law*, 1 LEGAL STUD. 267, 268 (1981) ("Both the leading textbooks [of English law] describe recklessness as deliberate risk-taking.").

his negligent failure to inquire.<sup>75</sup> In this chapter, I will not defend any particular position in this dispute. But though my arguments leave certain questions unanswered, they resolve the conceptual issues that are distinctive to negligence. By identifying the distinctive kind of risk that negligent agents create, I distinguish negligence from other grades of culpability, and by showing that the same mental state involved in other grades of mens rea may also explain culpability for negligence, I disarm what has commonly been seen as the primary obstacle to culpability for negligence—namely, showing how individuals may be culpable for creating a risk they do not perceive. The problem of identifying the mental state required for negligence does remain. But that problem is not distinctive to negligence, because it must be confronted in any event to understand other grades of mens rea themselves.

### *B. Distinguishing Negligence*

Of course, this project succeeds only if my account does successfully distinguish negligence from other grades of mens rea. But given my analysis of negligence it may seem unclear whether I have, in fact, characterized recklessness and negligence as different kinds of culpability at all. By understanding recklessness and negligence each to be a kind of culpable disregard of risk, after all, my account in some sense assimilates the two. In what sense, then, are they distinct grades of culpability, on my view? They do involve the culpable disregard of different risks. But not all acts of recklessness must in general involve the same risks: individuals who consciously disregard a wide range of different risks may nonetheless all count as reckless, so long as they grossly deviate from the standard of conduct of the law-abiding person. The differences among those risks do not correspond to different kinds of culpability; rather, those are simply different risks that may all be disregarded recklessly. Even if negligence does involve the disregard of a different sort of risk than recklessness, then, why should that difference distinguish negligence as a sort of culpability distinct from recklessness? Why is the conscious disregard of the risk that omitting inquiry will prevent individuals from accurately estimating the

---

<sup>75</sup> *E.g.*, Kenneth W. Simons, *Rethinking Mental States*, 72 B.U. L. REV. 463, 486–90 (1992) (articulating a conception of recklessness as indifference).

risks their actions create different from the conscious disregard of any other sort of risk?

The basis I propose to distinguish recklessness from negligence is the kind of risk that each defendant disregards: the reckless defendant simply disregards the risk that his action will cause harm, while the negligent disregards the risk that his failure to inquire will prevent him from identifying evidence that his action will create an unreasonable risk of harm. Whether this difference can adequately explain the distinction between recklessness and negligence, in turn, depends on what the function of that distinction is. The law does not separate the two grades of mens rea merely for taxonomic clarity. Instead, it treats negligence as a less serious form of culpability than recklessness: offenses defined with the mens rea of negligence are typically of a lower degree than otherwise identical offenses defined with recklessness,<sup>76</sup> and many crimes defined using recklessness would not be criminal if performed only negligently.<sup>77</sup> Any account of the difference between negligence and recklessness, then, must explain why that distinction can justify reducing the severity of negligent defendants' sentences or even exempting them from punishment entirely. The distinction I have identified between recklessness and negligence provides such an explanation: because of structural differences in how I characterize negligence and recklessness, negligence systematically creates less risk than recklessness. Although negligence and recklessness do each involve the culpable disregard of a risk, then, different sentences are justified for each because of these differences in the degrees of risk they create.

On my account of negligence, negligent and reckless defendants are each culpable in virtue of the same kind of mental state: each culpably disregards a risk created by their conduct. But although the subjective aspects of their culpability are therefore the same, the objective aspects are not: the risk that each defendant disregards arises from a different aspect of the action he performs. And these differences in how each sort of action creates risk lead to differences in how risky

---

<sup>76</sup> Compare MODEL PENAL CODE § 210.3 (defining manslaughter, the reckless commission of criminal homicide, as a felony of the second degree) with *id.* § 210.4 (defining negligent homicide as a felony of the third degree).

<sup>77</sup> *E.g.*, MODEL PENAL CODE § 211.1(1)(a)–(b) (defining simple assault to include purposely, knowingly, or recklessly causing bodily injury, but not causing it negligently, unless done with a deadly weapon).



each action is. The Model Penal Code defines recklessness in terms of a “standard of conduct”; the defendant must have deviated from that standard, it specifies, in his disregard of risk.<sup>78</sup> As the name indicates, a standard of conduct governs conduct: it specifies that certain actions may not be performed. Thus, a defendant is reckless because of how dangerous the action he performed was, given the information about it he possessed. In particular, it was risky enough that he should not have performed it, “considering the nature and purpose of the actor’s conduct and the circumstances known to him.”<sup>79</sup> When a defendant is reckless, the risk that he disregards must reach this level: it must be a risk of a high enough degree for acting to be prohibited.

On my analysis of negligence, though, the degree of risk that a negligent wrongdoer disregards is lower, *ceteris paribus*, than the degree of risk disregarded by a similarly situated reckless wrongdoer who is aware of the evidence of danger that a negligent agent does not perceive. Because negligent agents fail to inquire before acting, they possess less information about their action than reckless agents who do uncover evidence through inquiry. They differ in the evidence they possess about their action: while a reckless defendant is aware of the substantial and unjustifiable risk that it creates, a negligent defendant who fails to inquire is unaware of that risk because he lacks the evidence of danger that inquiry would have revealed. From the negligent defendant’s perspective, inquiry is not certain to reveal any such evidence. Before a driver turns to see if a car is coming, for example, there is a chance that the road will be clear. Thus, possibilities in which the road is clear must be employed in computing the degree of risk “considering the nature and purpose of [the defendant’s] conduct and the circumstances known to him.”<sup>80</sup> Furthermore, harm is obviously less likely to occur in possibilities where no additional evidence of danger exists than in possibilities where it does exist. Since negligent agents who fail to inquire must include both sorts of possibilities in evaluating risk, while reckless agents must include only those in which the additional evidence does exist, the degree of risk disregarded by a negligent defendant will be lower than the risk disregarded by a reckless one. In particular, a negligent agent disregards a risk equal in degree to the

---

<sup>78</sup> MODEL PENAL CODE § 2.02(2)(c).

<sup>79</sup> MODEL PENAL CODE § 2.02(2)(c).

<sup>80</sup> MODEL PENAL CODE § 2.02(2)(d).

risk a reckless agent disregards, further discounted by the probability that inquiry will yield evidence of danger. Reckless defendants disregard a risk, while negligent defendants merely disregard the risk of a risk—the risk that inquiry would reveal their action to create a substantial and unjustifiable risk. And this additional iteration of risk reduces the likelihood of harm.<sup>81</sup>

Of course, even without further inquiry the evidence a defendant already possesses about his action may be sufficient for him recognize that it creates a substantial and unjustifiable risk. In those cases, the action is simply reckless. For the reasons I have given, the risks those defendants consciously disregard may be lower than the risks that would be disregarded by similarly situated defendants who do inquire before acting and thereby gain additional evidence of danger. But so long as the risk an action creates is sufficiently high that performing that action would grossly deviate from the standard of conduct of the law-abiding person, the defendant who consciously disregards the risk is reckless. Because risks of that degree are sufficiently high that actions creating them are prohibited, someone who recognizes a possible action to create a risk of that degree simply should not perform it. It is reckless to disregard a risk of a degree high enough that the action creating it simply should not be performed. The fact that further inquiry might have revealed the action to be even more risky is simply beside the point.<sup>82</sup>

Obviously, when defendants fail to inquire in contexts such as these, the degree of risk they disregard is not lower than the degree of risk ordinarily created by reckless defendants, because the failure to inquire is not the primary source of the risk their actions create. Instead, because such defendants are already aware of the substantial and unjustifiable risk their actions will create, given the evidence they already possess, they consciously disregard a risk of the same degree as other reckless

---

<sup>81</sup> This inequality can be expressed formally. Let  $p_1$  be the likelihood that an action will cause harm  $H$ , given the information available to a reckless agent. The reckless agent thus disregards a risk of magnitude  $p_1H$ —that is, the harm multiplied by the likelihood that it will occur. The negligent agent, however, is unaware of certain evidence available to the reckless agent. Let  $i$  be the likelihood that inquiry will reveal that evidence, and let  $p_2$  be the much lower likelihood of harm in possibilities where inquiry would reveal that no such evidence exists. The negligent agent disregards a risk of magnitude  $i(p_1H) + (1 - i)(p_2H)$ —the likelihood of harm given evidence multiplied by the likelihood of evidence plus the likelihood of harm without evidence multiplied by the likelihood of no evidence. And, since  $p_2$  is less than  $p_1$ , by stipulation, this magnitude is lower than  $p_1H$ .

<sup>82</sup> Whether such defendants failed to inquire may be relevant to whether their culpability should be elevated from recklessness to knowledge, of course. See *supra* notes 54–68 and accompanying text.

defendants. But obviously not every agent who fails to inquire will be aware of risks of this degree. Sometimes, prior to further inquiry the risk that an action creates, evaluated from the agent's perspective, will not be sufficiently high to forbid performing that action. Such risks do not rise to the level at which disregard would grossly deviate from the standard of conduct of the law-abiding person. Considered on their own, then, such actions are not so risky that performing them is simply prohibited as reckless. These actions, I will argue, may sometimes be negligent. Indeed, common law definitions of involuntary manslaughter, the least culpable form of homicide, explicitly specify that the negligent defendant commits a crime even though his act is not criminal in itself: as the Michigan Supreme Court explains, for example, involuntary manslaughter (in one of its forms) consists of "negligently doing some act lawful in itself."<sup>83</sup> Such defendants are negligent, but they are not reckless: given their information about the action they perform, disregarding the risk it creates does not grossly deviate from the standard of conduct of the law-abiding person.

Negligence is a lower grade of culpability than recklessness: defendants who satisfy the definition of recklessness automatically count as negligent.<sup>84</sup> Thus, if a defendant recognizes a risk substantial enough for him to count as reckless, the question of whether he satisfies the definition of negligence need not arise. Instead, a theory of negligence must describe a kind of culpability that applies to individuals who do not satisfy the definition of any higher grade of mens rea. Negligent defendants, then, do not recognize a risk of a high enough degree that their action is prohibited. Nonetheless, they fail in a different way to comply with the requirements imposed by law: they grossly deviate from a "standard of care" concerning the perception of risk.<sup>85</sup> A standard of care does not govern conduct: it does not forbid certain actions as impermissible in themselves. Rather, it governs care—the steps individuals take to ensure that their conduct does not violate the rules. Thus, although

---

<sup>83</sup> *People v. Holtschlag*, 684 N.W.2d 730, 739 (Mich. 2004) (citing *People v. Ryczek*, 194 N.W. 609, 611 (Mich. 1923)); *see also* *State v. Thomas*, 211 A.3d 274, 285 (Md. 2019) ("There are generally thought to be three varieties of involuntary manslaughter . . . (2) gross negligence manslaughter—'negligently doing some act lawful in itself' (citing *State v. Albrecht*, 649 A.2d 336, 499 (Md. 1994)). The phrase dates to Blackstone. *See* 4 WILLIAM BLACKSTONE, COMMENTARIES \*192.

<sup>84</sup> MODEL PENAL CODE § 2.02(5).

<sup>85</sup> MODEL PENAL CODE § 2.02(2)(d).

given the information the negligent agent possesses about his action it is not so risky that it must not be performed, his decision not to inquire before performing it breaches a separate standard governing the care that must be taken in inquiring before acting. Although the risk he disregards is not so high to require him to refrain from acting entirely, it is high enough to require him to inquire before performing it. Negligent and reckless agents both do consciously disregard a risk in acting, but in disregarding risk they each breach a different kind of standard governing their conduct. And the thresholds at which risks breach each standard are different. The reckless agent disregards a risk substantial enough to breach a standard specifying when actions are too risky to be performed. The negligent agent, by contrast, disregards a lower risk, which is not high enough to violate a standard governing which actions may be performed. But even though the risk is lower, it does violate a different standard that specifies when individuals must inquire before acting.

Schematically, then, recklessness and negligence differ in the structures of the decisions they govern. When agents act recklessly, they face a choice between only two normatively significant options—namely, whether or not to perform a particular action. The standard of conduct imposed by prohibitions on recklessness governs that decision: if, given the information the agent possesses about an action, the risk it creates is substantial and unjustifiable, he may not perform it. When agents act negligently, by contrast, they face a decision with a more complex structure. Negligent agents have three options—to act, to refrain from acting entirely, or to inquire further before acting. The standard of care imposed by prohibitions on negligence would not dictate any choice between the options of acting and refraining from action on their own: the degree of risk the action creates is not so high that acting would in itself grossly deviate from the standard of conduct of the law-abiding person. But a third option is present: the agent could defer choosing how to act until after he inquires. And the standard of care imposed by prohibitions on negligence governs the choice between those three options: in particular, it requires individuals to choose inquiry before acting over acting alone. The negligent agent does not perform an action that is so risky, taken in itself, that no law-abiding person would perform it, but the manner in which he performs that action—in particular, his

omission of inquiry before performing it—is risky enough to fall short of the care that a reasonable person would exercise.

On my account, then, though reckless and negligent agents do each violate a legal prohibition by consciously disregarding a risk in acting, they violate those prohibitions in virtue of different aspects of their conduct. Reckless agents are culpable for performing an action that creates an unacceptably high risk of harm. By prohibiting recklessness, then, the law requires agents simply not to act: if an action is too dangerous to perform, given the information they possess about it, they must simply refrain from performing it. By contrast, negligent agents are culpable for not having inquired before they acted. Thus, while the law requires reckless defendants not to perform their reckless action at all, it does not generally require negligent agents not to perform their negligent action. Instead, it requires them only to employ a particular precaution before performing it—namely, to engage in inquiry first. Had they first taken that precaution, they could have gone on to act—at least, unless inquiry uncovered evidence that acting would create a substantial and unjustifiable risk.

Reckless and negligent agents are alike, then, in the kind of mental state responsible for their culpability. But although in each case acting with that mental state is prohibited, those prohibitions guide individuals' conduct in different ways. Prohibitions on recklessness instruct individuals not to perform actions that are substantially and unjustifiably risky in light of the evidence available about them. Prohibitions on negligence, instead, instruct individuals that they must sometimes take a particular precaution before acting—namely, inquiring—if there is a high enough risk that omitting that precaution will increase the risk their action creates by preventing them from acquiring evidence that would lead them to modify their conduct. Defendants can be negligent without being reckless, then, because these two prohibitions employ different thresholds: a risk can be high enough that individuals must inquire before acting but low enough that they are not prohibited from acting entirely.

According to this analysis, then, an action can violate the standard governing when the disregard of risk in acting would be reckless without violating the standard

governing when the omission of inquiry prior to acting would be negligent. Reckless and negligent agents are each culpable for their conscious disregard of a risk: each, that is, performs a particular action in disregard of the degree of risk that action will create. Thus, determining whether an individual was either reckless or negligent involves evaluating the same mental state—his disregard of a risk of some particular degree in acting. By suggesting that disregarding risk may be negligent without being reckless, I propose that these two standards can disagree about whether disregarding a risk of the very same degree is culpable. But it is unclear why the degree of risk that the criminal law permits an individual to disregard would differ based on which provision of law he consults. Instead, one might think, there is simply a particular threshold above which the law prohibits individuals from imposing risks on others. On this approach, there would seem to be no room for negligence, as I construe it: the conscious disregard of risks above that threshold would be reckless, and the conscious disregard of risks below it would not be culpable. Why, then, would the law instead impose different standards governing whether risks may be disregarded that disagree about whether a particular action was too risky? If an action isn't risky enough to be reckless, why isn't it simply not risky enough to be a crime?

Prohibitions on recklessness and negligence may employ different thresholds to determine when risk-creation is permissible, though, because those prohibitions govern different kinds of decisions. Recklessness, as I have analyzed it, governs the choice between whether or not to perform a particular action. By contrast, negligence governs the choice between whether to perform a particular action or to inquire first, then perform it unless further evidence reveals it to be reckless. These choices involve different alternatives: they concern whether the action creates an acceptable degree of risk as compared to two different courses of conduct that the defendant might pursue instead. And, in general, whether a particular choice is acceptable often depends on what alternatives to that choice are available. The reasons for and against performing a particular action affect the outcome of any comparison between that action and another, but that comparison also depends on the reasons for and against the alternative. And different reasons will obviously weigh for and against different alternatives. Thus, recklessness and negligence can criminalize creating risks of

different degree due to differences in the reasons for or against the alternative course of conduct that each requires the defendant to pursue.

Whether actions may be negligent even though they are not risky enough to be reckless, then, depends on the differences between the courses of conduct that prohibitions on recklessness and negligence require individuals to pursue. On my view, recklessness imposes standards of conduct that identify when individuals may not perform actions, while negligence imposes standards of care that identify when individuals must take a particular precaution—inquiring before acting. And obvious differences between inquiring before acting and avoiding an action entirely can lead to differences in comparing a single action that creates a particular degree of risk to each of these different alternatives. The costs and benefits involved in a comparison between acting and refraining from acting are relatively straightforward. Acting has certain costs: it creates a risk of harm. At the same time, though, it has benefits—whichever goods the agent aims to achieve by performing it. Avoiding the action avoids both its costs and its benefits. Thus, although individuals can avoid creating a particular risk by avoiding the action that creates it, they must also sacrifice the benefits the action would bring. Whether they should perform or avoid it, then, simply depends on whether the costs outweigh the benefits.

While requiring individuals not to act avoids the costs of acting but also avoiding the benefits, requiring inquiry as a precaution reduces those costs more efficiently. Since standards of reasonable care would not require inquiry unless it will likely reveal evidence affecting the choice the agent should make, any evidence that required inquiry reveals is likely to substantially affect how an action's risk should be estimated. Drivers should look before turning, for example, because the danger of turning is dramatically affected by the presence of oncoming traffic. Inquiry reduces risk because it enables individuals to avoid acting in the subset of possibilities in which acting is very dangerous: an individual who inquires will not act if he uncovers evidence of danger and will act only if he does not. By allowing individuals to avoid acting in those possibilities, inquiry can eliminate most of the expected costs of acting: since inquiry allows an individual to avoid acting when harm is much likelier to occur, but not when it is much less likely to occur, the reduction in expected harm is

disproportionate to the likelihood that the evidence actually exists. But because the benefits of acting are likely to be the same whether or not that evidence exists, inquiry reduces the expected benefits of acting by much less.<sup>86</sup> Individuals cannot capture those benefits if the additional evidence uncovered prevents them from acting, but since the existence of that evidence makes no difference to the magnitude of those benefits, the reduction in expected benefits is simply the likelihood that inquiry will uncover that evidence. The reduction in expected benefits is therefore lower than the reduction in expected costs.

While simply prohibiting an action eliminates its risks at the cost of eliminating its benefits, then, inquiry instead produces a reduction in risk that is disproportionately large compared to the reduction in expected benefits. Consequently, even if an action does not create a risk of so high a degree that it simply should not be performed, sometimes individuals should nonetheless inquire before they perform it. If its risks do not exceed its benefits, avoiding it would prevent more good than harm and thus would be unjustified. But inquiry may sometimes be justified because it can reduce most of the risk an action creates without also reducing most of its benefits, for even if the total benefits do exceed the total harms, a large fraction of the harm may nonetheless exceed a small fraction of the benefit. In addition, of course, inquiry will often incur costs of its own. When there is no easy way to inquire before acting, those costs may be high enough that, on the whole, inquiry is not required. But when inquiry is easy, as it sometimes will be, its total benefits will exceed its total costs, even though the action's own risks are not so great as to require avoiding it. In such cases, the choice of whether or not to perform an action should be answered differently than the choice of whether to inquire before performing it.

Thus, the standards governing when actions are prohibited may coherently differ from the standards governing when inquiry is required: an action may be risky enough that individuals should inquire before acting even though it is not risky enough that they should refrain from performing it entirely. As a result, an action that is not risky enough to be reckless nonetheless may be risky enough to be negligent.

---

<sup>86</sup> Indeed, inquiry may sometimes provide benefits, too, if the agent risks suffering harm himself, too.



Criminal prohibitions prevent individuals from performing actions that substantially threaten various important interests that the law aims to prevent. But such prohibitions come at a cost: they restrict individuals' autonomy by dictating the choices they must make.<sup>87</sup> Thus, the law should not restrict individuals' autonomy by prohibiting them from freely choosing to perform certain actions unless that restriction is justified by the threat those actions pose to the interests that criminal law aims to protect. As the risk that an action will actually cause a particular kind of harm decreases, of course, the threat it poses decreases as well, since the actions that fail to cause harm do not ordinarily invade the interests the law aims to protect. But a prohibition on acting obviously continues to restrict autonomy to the same extent even as risk declines. At some point, then, that risk will decrease to the point that it can no longer justify the restriction in autonomy that accompanies criminal prohibitions. Once risks fall below that point, consciously disregarding them is no longer reckless.

While a decrease in the degree of risk that an action creates may prevent the threat that it poses from justifying the restrictions on autonomy associated with prohibiting it, even that lower degree of risk might be threatening enough to justify a criminal prohibition that imposes less severe restrictions on the autonomy of those it binds. And, as I have argued, the requirement to inquire before acting constitutes a less burdensome restriction on autonomy than a prohibition on acting entirely. While a prohibition always frustrates the preferences of those who wish to act, because it always requires them not to act, the requirement to inquire usually allows individuals to act as they wish, since they need only alter their conduct in the relatively few instances when inquiry actually does produce evidence of danger. Because individuals' autonomy is more seriously restricted by the requirement always to refrain from performing a particular action than by the requirement only to inquire, then sometimes to refrain from acting, the latter restriction may be justified even if risks are insufficiently threatening to justify the former. Some actions that are not

---

<sup>87</sup> For a classic statement of the view that criminal law's joint aims are to preserve individual freedom while restricting harmful conduct, see H.L.A. Hart, *Prolegomenon to the Principles of Punishment*, in PUNISHMENT AND RESPONSIBILITY: ESSAYS IN THE PHILOSOPHY OF LAW 1, 23 (2d ed. 2008) (characterizing criminal law as "a method of social control which maximizes individual freedom within the coercive framework of law").

sufficiently risky to be prohibited as reckless may nonetheless be sufficiently risky that the failure to inquire before performing them may be prohibited as negligent.

Ultimately, risk-creation cannot be eliminated entirely from human affairs; though we should attempt to avoid unnecessary risks, our imperfect knowledge of the future makes inevitable some degree of unintended harm. Because we cannot act at all without creating some degree of risk, we must be able to act when the risk is sufficiently low: though of course there are difficult questions about where that threshold lies, at some point as a society we value our ability to engage in an activity more greatly than we value the harms we might avoid by foregoing that activity. We are generally justified in acting when we create risks that lie beneath that threshold, because those risks are necessarily incurred by activities we are unwilling to forego. But because risks that exceed that threshold are not justified as the inevitable byproducts of such activities, they are prohibited; the law guides individuals to avoid them by defining their disregard as culpable recklessness. Just because risks beneath that threshold are justified if incurred necessarily, though, it does not follow that all risks beneath that threshold must be justified, for not all risks beneath it are incurred necessarily. In particular, some actions create risks beneath that threshold unnecessarily because by inquiring before acting the agent could have reduced them considerably. And since inquiry does not generally prevent individuals ultimately from acting, these risks are not a necessary byproduct of activities we are unwilling to forego. Consequently, the justification we possess in other cases to create risks of that degree does not apply. The law guides individuals to avoid these risks, then, by defining their disregard as culpable negligence. Because they could have been avoided by inquiry, we are not justified in creating them, even though risks of that degree might be justified were it possible to avoid them only by abandoning the activities that create them.

As I have noted, the criminal law distinguishes between recklessness and negligence in grading them differently: negligence is a less culpable form of mens rea than recklessness. An account of negligence, then, must explain why the law would treat it as less serious than recklessness. On my view, negligence is a distinctive form of culpable wrongdoing. While agents act recklessly by disregarding a substantial and

unjustifiable risk created by the action they perform, agents act negligently by omitting a particular precaution—inquiry—in disregard of the increase in the action’s risk that inquiry would avoid. Because restricting actions imposes more substantially on individuals’ freedom of choice than requiring inquiry does, the failure to inquire before acting may justifiably be prohibited as negligent in contexts where the action is not itself risky enough to justify prohibiting its performance as reckless. Thus, negligent wrongdoing systematically creates a lower degree of risk than reckless wrongdoing does. This difference, in my view, explains why negligence is a less serious form of mens rea than recklessness.

In general, the mens rea distinctions drawn by the Model Penal Code are sensitive to the degree of risk that an action creates. This sensitivity is most obvious in the distinction between knowledge and recklessness. Under the Model Penal Code, a person commits an offense knowingly when he is aware that the elements of that offense are satisfied or are “practically certain” to be satisfied, while he acts only recklessly when he is conscious of a risk that those elements are or will be satisfied.<sup>88</sup> Thus, while the knowing agents disregard the maximum degree of risk of performing the elements of an offense, reckless agents disregard a lower degree of risk—a “probability less than substantial certainty,” as the commentaries put it.<sup>89</sup> Knowledge, of course, is a more serious form of culpability than recklessness.<sup>90</sup> The reduced risk of recklessness compared to knowledge, then, corresponds to a reduction in the seriousness of reckless wrongdoing compared to knowing wrongdoing. My account of negligence simply extends this principle as a general explanation of the relative seriousness of the grades of mens rea defined by the Model Penal Code.<sup>91</sup> Just as recklessness involves less risk than knowledge, negligence involves less risk than recklessness; thus, just as recklessness is a less serious kind of mens rea than knowledge, negligence is a less serious kind of mens rea than recklessness. Consequently, the basis I identify for distinguishing negligence from recklessness can

---

<sup>88</sup> Compare MODEL PENAL CODE § 2.02(2)(b) (AM. LAW INST. 1962) with *id.* § 2.02(2)(c).

<sup>89</sup> *Id.* § 2.02 cmt. 3.

<sup>90</sup> *Id.* § 2.02(5).

<sup>91</sup> My proposal admittedly does not explain the distinction between purpose and knowledge, which the Model Penal Code’s commentaries themselves acknowledge to be “inconsequential for most purposes of liability.” *Id.* § 2.02 cmt. 2.

plausibly explain how negligence relates to recklessness in the Model Penal Code's mens rea hierarchy.

Of course, this appeal to a general principle correlating an offense's mens rea with the magnitude of the risk it disregards might not be plausible if that general principle were itself normatively unsound. The Model Penal Code's distinctions between forms of mens rea affect both the scope and the magnitude of criminal liability: if an action satisfies the elements of an offense, the agent's mens rea may affect both whether and how severely his action will be punished.<sup>92</sup> Thus, whether my account can explain how negligence fits into the mens rea hierarchy depends on whether the degree of risk an action creates should affect both whether and how severely it is punished. While this claim does strike me as intuitively plausible, it obviously depends on more general normative principles governing the scope and degree of justified punishment. I will not defend any account of those principles here. But I will argue that allowing the degree of risk an action creates to influence whether and how severely it is punished advances at least two purposes that punishment is often understood to serve—namely, retribution and deterrence—which at least *prima facie* supports understanding the mens rea hierarchy through the degrees of risk that different forms of mens rea create.<sup>93</sup>

Retributivist accounts of punishment broadly take punishment to be a deserved response to wrongdoing.<sup>94</sup> Little consensus exists among retributivists about what, exactly, determines the extent to which any particular action deserves to be punished.<sup>95</sup> But since punishment is a response to the wrongful features of the action, in general punishment must be proportional to the gravity of the wrong for which it is

---

<sup>92</sup> See *supra* notes 76–77 and accompanying text.

<sup>93</sup> Obviously, rival accounts of the purpose of the purpose of punishment also exist, including the incapacitation or rehabilitation of offenders. See, e.g., 18 U.S.C. § 3553(a) (2018) (“The court, in determining the particular sentence to be imposed, shall consider . . . (2) the need for the sentence imposed . . . (C) to protect the public from further crimes of the defendant; and (D) to provide the defendant with needed educational or vocational training, medical care, or other correctional treatment in the most effective manner . . .”). I do not mean to take a position on whether these purposes might also justify sentencing offenders less severely for disregarding a less substantial risk.

<sup>94</sup> See Zachary Hoskins & Antony Duff, *Legal Punishment*, STAN. ENCYCLOPEDIA PHIL. § 4 (Edward N. Zalta ed., Summer 2022 ed.), <https://plato.stanford.edu/entries/legal-punishment/>; Alec Walen, *Retributive Justice*, STAN. ENCYCLOPEDIA PHIL. (Edward N. Zalta ed., Summer 2022 ed.), <https://plato.stanford.edu/entries/justice-retributive/>.

<sup>95</sup> Indeed, as Yaffe notes, “[i]t is striking how rarely philosophers have made an effort to offer a theory of desert, given how much has been written about the concept.” GIDEON YAFFE, *THE AGE OF CULPABILITY: CHILDREN AND THE NATURE OF CRIMINAL RESPONSIBILITY* 99 (2018).

imposed.<sup>96</sup> The risk an action creates may affect the scope and magnitude of justified punishment, then, just in case it affects how wrongful the action is. And while different accounts of the nature of wrongs are obviously possible, riskier actions seem clearly to be more wrongful, since they pose a greater threat to the interests that they risk harming. Thus, riskier actions plausibly deserve greater punishment than less risky ones.<sup>97</sup> This difference in desert, in turn, justifies how different kinds of mens rea affect sentences.

Since disregarding lower degrees of risk is less wrongful, offenses performed with a mens rea involving less risk should be sentenced less severely. To be sure, this rule holds only *ceteris paribus*; the degree of risk an action creates is only one factor affecting how wrongful it is, and thus every reckless offense certainly should not be sentenced more severely than every negligent one. But sentencing law obviously does not sentence offenders based solely on their mens rea; instead, a wide range of additional factors—including, most obviously, the elements of the offense of conviction—also affect the sentence imposed.<sup>98</sup> Because different sentencing systems employ different methods of integrating mens rea with other factors, those systems can only be evaluated individually, not in the abstract. The general approach that they

---

<sup>96</sup> Walen, *supra* note 94, at 4.4.

<sup>97</sup> In part because there is little consensus about exactly how an action's wrongfulness affects its desert (and because establishing such a consensus is not among this chapter's ambitions), I will not defend this claim by developing a detailed analysis of how desert affects wrongdoing. But that the greater desert of riskier offenses follows naturally from at least one recent effort to spell out how an action's characteristics affect the degree of punishment it deserves—Yaffe's view that desert depends on "the gap between (1) the reason-giving weight one is disposed to grant to a given fact in one's deliberations, and (2) the weight of the legal reason provided by the fact." Gideon Yaffe, *Legal Reasons, Legal Desert, Legal Culpability: Reply to Guerrero, Kelly and Mendlow*, 24 J. ETHICS 295, 295 (2020). The greater a risk, the weightier the legal reason against disregarding it, since the expected harm of disregarding it is greater. Thus, *ceteris paribus*, someone who disregards a risk of a higher degree is more deserving of punishment: the gap between the correct weight of that reason and the inadequate weight an agent assigns it will ordinarily be greater the higher the correct weight is.

<sup>98</sup> For example, negligent homicide is a third-degree felony while reckless simple assault (or indeed any other kind of simple assault) is a misdemeanor. Compare MODEL PENAL CODE § 210.4 (AM. LAW INST. 1962) with *id.* § 211.1. In addition, the Model Penal Code's sentencing provisions generally allow a wide range of other factors to affect the sentence ultimately imposed. As initially drafted, it merely limited the upper and lower bounds of the indeterminate sentences that were authorized, thereby allowing judges to consider a wide range of factors in determining the sentence actually imposed. *E.g.*, *id.* § 6.06. The sentence within that range actually served itself depended on the decision of the relevant Board of Parole, which itself could consider a wide range of very different factors in deciding whether to release a prisoner. See *id.* § 305.9–.10. More modern approaches to sentencing, such as federal sentencing law or the Model Penal Code's revised sentencing provisions, replace indeterminate sentences with complex rules that precisely define a narrow sentencing range for every offense based on a detailed consideration of features of both offense and offender. See generally U.S. SENTENCING GUIDELINES MANUAL (U.S. SENTENCING COMM'N 2021); MODEL PENAL CODE: SENTENCING art. 6B (AM. LAW INST., Proposed Final Draft 2017). Mens rea distinctions are only one of many factors that such guidelines typically consider.

take, though, is consistent with the role that mens rea distinctions would play in sentencing were they to track the degree of risk an offender disregards in acting: risk should influence sentences to some extent without completely determining them, just as mens rea distinctions do under existing law.

In addition to sometimes sentencing negligent offenders less severely than reckless ones, the law also often exempts negligent wrongdoers from criminal punishment entirely, by defining many offenses that cannot be performed negligently.<sup>99</sup> Such exemptions might seem puzzling, on my account. Even if negligence is less risky than recklessness, and thus less deserving of punishment, it still deserves some degree of punishment. Why, then, would it not always be punished to some extent, just less severely than recklessness? This puzzle confronts any version of retributivism that accepts culpability for negligence, though: if negligence is culpable, then some degree of punishment would seem justified for negligently causing any sort of harm that is punishable if caused recklessly, even though the law punishes very few negligent offenses.<sup>100</sup> Since retributivism holds culpable wrongdoing to deserve punishment, considerations beside desert must explain the exemption of most culpable negligence from punishment. But it is fairly easy to see what kinds of considerations might justify this restriction on liability. Even if punishing the guilty is intrinsically good, it is not the only good: retributivists can accept the importance of desert without insisting that the state must always give offenders what they deserve, no matter what the cost.<sup>101</sup> If an individual deserves very little punishment, the good achieved by administering it simply may not be enough to justify invoking the machinery of criminal law. And because actions that create little risk deserve little punishment, imposing punishment on such agents will often not be worth it. Because the degree of risk an action creates is not the only factor affecting desert, of course, sometimes negligence should still be punished—for example,

---

<sup>99</sup> Only the following Model Penal Code crimes may be committed negligently: criminal homicide, MODEL PENAL CODE § 210.4, assault, but only if caused with a deadly weapon, *id.* § 220.3(1)(a), and criminal mischief, but only if caused in the employment of fire, explosives, or other dangerous means, *id.* § 211.1(b).

<sup>100</sup> The existence of some degree of desert would seem to be a direct consequence of an analysis of desert in terms of harm and culpability: if negligence is a form of culpability, and if it causes harm, then punishment is deserved. Alec Walen, *Supplement to Retributive Justice*, STAN. ENCYCLOPEDIA PHIL. 1 (Edward N. Zalta ed., Summer 2022 ed.).

<sup>101</sup> For a defense of this claim, see Douglas Husak, *Holistic Retributivism*, 88 CALIF. L. REV. 991 (2000).

negligently causing a very serious harm like death. But disregarding a low risk of a less serious harm may not deserve enough punishment to justify actually imposing it. For that reason, negligence often should not be punished, even when it may deserve some punishment.

Accounts of punishment based on deterrence argue that punishing offenders may increase overall welfare because such punishments communicate a threat that induces potential offenders to refrain from offending.<sup>102</sup> The details of effective deterrence are complex and contested. But, in general, since punishment is justified by the harm it prevents, the amount of harm caused by a particular type of action affects how actions of that type may justifiably be punished. Since negligent actions are less risky than reckless ones, they cause less harm in total; thus, sometimes deterrence will justify less severe sentences for negligence, and sometimes it will not justify punishment at all. First, punishment is justified only if the harms that would have been produced by the offenses deterred exceed the harms produced by the imposition of punishment itself, plus other associated costs.<sup>103</sup> Thus, the amount of punishment that may justifiably be imposed to deter a particular kind of offense is limited by the amount of harm offenses of that kind cause.<sup>104</sup> Since negligence is less risky than recklessness, negligent wrongdoing generally causes less harm than reckless wrongdoing. Thus, less total punishment may be imposed for a particular type of negligent wrongdoing than for the analogous type of reckless wrongdoing. And, of course, that amount of punishment may be reduced by sentencing offenders less severely for negligence than for recklessness.<sup>105</sup> Furthermore, sometimes the total amount of harm created by negligent conduct may be so small that it is outweighed

---

<sup>102</sup> *E.g.*, CESARE BECCARIA, ON CRIMES AND PUNISHMENTS AND OTHER WRITINGS 26 (Aaron Thomas ed., Aaron Thomas & Jeremy Parzen trans., 2008); JEREMY BENTHAM, AN INTRODUCTION TO THE PRINCIPLES OF MORALS AND LEGISLATION 157 (J.H. Burns & H.L.A. Hart eds., 1996).

<sup>103</sup> BENTHAM, *supra* note 102, at 159 (specifying that punishment should not be imposed “where the mischief it would produce would be greater than what it prevented”).

<sup>104</sup> *See id.* at 168 (“The greater the mischief of the offence, the greater is the expence, which it may be worth while to be at, in the way of punishment.”).

<sup>105</sup> To be sure, other factors also limit the total amount of punishment imposed for negligence: at least when offenses are defined to require that a harmful result be caused, punishment would be imposed less often for negligence in exact proportion to the lower rate at which it causes harm. How this reduction in total punishment and a reduction in the sentences imposed jointly affect the effectiveness of deterrence is complex, and analyzing it lies beyond the scope of my argument here.

by the costs of deterring it through punishment.<sup>106</sup> In these cases, offenses should be defined so that they cannot be performed negligently. Thus, deterrence too can explain why the degree of risk an action creates should influence punishment in a manner consistent with how I have understood the distinction between recklessness and negligence.

### C. *The Scope of Criminal Negligence*

My account explains the culpability of criminal negligence, then, by reconceptualizing what it is—not a distinctive form of culpability for creating a risk that an individual fails to perceive, but rather a familiar form of culpability for a distinctive kind of risk. Since negligence is epistemic recklessness, it is culpable for the same reasons that recklessness is. One objection to this approach, which I have already discussed, would argue that it separates negligence too little from recklessness. Another objection, though, would attack it from the opposite end, so to speak, arguing that it separates criminal negligence too much from other forms of negligence. In particular, because of how I assimilate negligence to recklessness my account does not address the criminal culpability of defendants whose negligence does not involve the conscious disregard of a risk, however that notion is to be understood. Many have thought that these cases—instances of “true” or “pure” negligence, as some have put it—are the core examples that a theory of negligence must explain.<sup>107</sup> If my account does not justify culpability in those cases, it may seem that I have simply changed the subject. To be sure, since recklessness is a genuine form of criminal culpability, the sort of recklessness that I equate with negligence may be a genuine form of culpability, too. But because my account does not explain the culpability of true negligence, it may seem inadequate as a theory of negligence.

This objection obviously presupposes that a theory of criminal negligence must be able to explain the culpability of true or pure negligence. But is not clear to

---

<sup>106</sup> BENTHAM, *supra* note 102, at 159 (arguing that punishment should not be imposed “where the mischief it would produce would be greater than what it prevented”).

<sup>107</sup> ALEXANDER, FERZAN, & MORSE, *supra* note 4, at 69 (“True negligence is inadvertent creation of unreasonable risks.”); Husak, *supra* note 19, at 210 (defining “true negligence” to cover “cases in which it does not occur to defendants that they have created a risk”); Gideon Rosen, *The Problem of Pure Negligence*, in AGENCY, NEGLIGENCE AND RESPONSIBILITY 15, 19 (George Pavlakos & Veronica Rodriguez-Blanco eds., 2021) (“Pure negligence is negligence in which there is no pertinent ill will anywhere in the causal background.”).



me why this expectation is reasonable. A theory of criminal negligence must successfully explain what it purports to be about—namely, the legal doctrine that prosecutors and courts employ to convict and sentence certain offenders.<sup>108</sup> But a theory of criminal negligence does not need to explain the culpability of individuals whom the law does not actually hold liable for criminal negligence. Thus, I must explain the culpability of true negligence only if criminal law actually punishes true negligence. And it seems to me there is little reason to think that it does. In general, convictions for offenses of negligence occur primarily in circumstances in which one would hardly expect individuals to be totally unaware of risk in the manner characteristic of true negligence. The list of examples in a leading criminal law treatise of convictions for involuntary manslaughter, an offense requiring negligence, all involve circumstances in which the risks of failing to inquire are obvious and well-known: the use of automobiles,<sup>109</sup> the use of firearms,<sup>110</sup> the failure to seek medical care,<sup>111</sup> the provision of medical care,<sup>112</sup> the delivery of dangerous drugs,<sup>113</sup> the operation of a business in an unsafe building,<sup>114</sup> leaving a baby in a dangerous environment,<sup>115</sup> blasting,<sup>116</sup> and railroad maintenance.<sup>117</sup> Of course, individuals who perform such activities carelessly will not in every case consciously disregard the risk of failing to inquire. Each of these activities is obviously dangerous, though, and in each inquiry can often substantially reduce risks. As Jeremy Horder observes, they generally involve “well-known and accepted safety standards” that are designed to ensure that individuals do not unreasonably threaten the interests of others.<sup>118</sup> When standards of this sort exist, it seems quite plausible that individuals are often aware of the risks of deviating from them. Thus, while perhaps my theory does not justify literally every conviction for criminal negligence, there seems to me little reason to

---

<sup>108</sup> I do not mean to deny that a theory can be revisionist, of course, just that it is reasonable to suspect a theory that requires large revisions to practice really to be theorizing about something other than what it claims to be about.

<sup>109</sup> WAYNE R. LAFAVE, CRIMINAL LAW 1054 n.182 (6th ed. 2017).

<sup>110</sup> *Id.*

<sup>111</sup> *Id.* at 1054 n.184.

<sup>112</sup> *Id.* at 1054 n.186.

<sup>113</sup> *Id.* at 1055 n.187.

<sup>114</sup> *Id.* at 1054 n.185.

<sup>115</sup> *Id.* at 1055 n.188.

<sup>116</sup> *Id.* at 1055 n.189.

<sup>117</sup> *Id.* at 1055 n.190.

<sup>118</sup> Horder, *supra* note 10, at 517–18.

think that truly negligent defendants typically are convicted of such offenses. In order to defend the scope of actual liability for criminal negligence, then, I may happily concede that true negligence is not criminally culpable.

To be fair, the account I present here does not yield an actual rule of law that could be applied to particular cases: negligence requires the same sort of conscious disregard as recklessness, but I have not analyzed that conscious disregard here.<sup>119</sup> Thus, whether my account can explain the actual extent of prosecutions and convictions for criminal negligence depends on what sort of psychological attitude is ultimately accepted as the basis of culpability for recklessness. In particular, if a very narrow conception of subjective awareness ultimately proves correct, then plausibly some actions that courts and juries are presently willing to treat as culpable should not be crimes.<sup>120</sup> But it is not clear to me that this consequence would constitute an objection to my theory of criminal negligence. For if the requirement of subjective awareness for criminal culpability is a good deal more demanding than most have supposed, there probably will be a good deal less criminal culpability than most have supposed, too, and the scope of criminal negligence should probably be narrower than it presently is. Indeed, if the received view of the scope of criminal liability is generally too broad, surely cases of true negligence will not be criminally culpable. In that case it would seem to be a strength, not a weakness, that my account would nonetheless justify some criminal liability for negligence even if the scope of culpability were on the whole restricted considerably.<sup>121</sup>

Of course, for my theory of criminal negligence not to aspire to explain true negligence does not express any broader legal or moral view about true negligence. Criminal punishment is not the only remedy that exists in law or morality: a negligent individual may be liable for compensatory damages—or even simply liable to

---

<sup>119</sup> Or, at least, I have not analyzed it in this chapter. See *supra* notes 73–75 and accompanying text.

<sup>120</sup> For example, Kimberly Kessler Ferzan defends an account of culpability on which individuals must articulate a reason linguistically (at least in their minds) in order to be culpable for disregarding it: “Without triggering our speech centers, we are not reasoning through our actions and are not appraising their moral worth. To hold us responsible in such cases would be to violate “ought” implies “can.”” Ferzan, *supra* note 74, at 638. Certainly, I am sure that defendants presently may be convicted of criminal negligence even when they do not linguistically articulate that failing to inquire would increase the risk of acting, but they could not be were an account of this sort of subjective awareness correct.

<sup>121</sup> Indeed, many advocates of relatively narrow theories of subjective awareness reject negligence liability entirely. ALEXANDER, FERZAN, & MORSE, *supra* note 4, at 69–85.

interpersonal blame—even though he is not liable to criminal punishment.<sup>122</sup> Not all bad acts are crimes. Accounts of criminal culpability that reject a requirement of subjective awareness are often motivated by reflection on cases of inadvertence that seem intuitively to justify condemning the agent: R.A. Duff, for example, analyzes the moral implications of inadvertence by discussing a groom who “missed his wedding” because “he was in the pub with his friends at the time, and . . . the wedding just slipped his mind.”<sup>123</sup> I do not mean to deny that this action would reflect quite poorly on the groom, and that it would justify responding negatively to him. (Most importantly—call off the wedding!) But such actions are not criminally punishable. Consequently, it does not seem to me that a theory of criminal negligence must explain their culpability. At the very least, those who would claim that a theory of criminal negligence must itself apply more broadly to forms of negligence that are not criminal should bear the burden of showing why.

One prominent understanding of the relationship between criminal negligence and ordinary negligence does perhaps explain why a theory of the former must explain the latter—namely, the analysis that criminal negligence is simply “gross negligence.”<sup>124</sup> On this approach, as Michael Moore and Heidi Hurd explain, the two differ only in degree: criminal negligence is “ordinary, civil negligence, only more of it.”<sup>125</sup> In a more memorable example of this approach, Calvert Magruder explained the difference between the two as the difference between a fool and a damn fool.<sup>126</sup> If criminal negligence is the same thing as ordinary negligence, just more, then a theory

---

<sup>122</sup> See, e.g., JOSHUA DRESSLER, UNDERSTANDING CRIMINAL LAW 127–28 (8th ed. 2018); Michael S. Moore & Heidi M. Hurd, *Punishing the Awkward, the Stupid, the Weak, and the Selfish: The Culpability of Negligence*, 5 CRIM. L. & PHIL. 147, 148–50 (2011); Simons, *supra* note 18, at 291–94.

<sup>123</sup> DUFF, *supra* note 10, at 163. Philosophers have been particularly fond of examples in which agents negligently forget facts that would make acting impermissible. E.g., ALEXANDER, FERZAN, & MORSE, *supra* note 4, at 77 (describing parents who forget a child in the bath); Husak, *supra* note 19, at 201 (describing a parent who forgets his infant in the bath); Smith, *supra* note 5, at 116–17 (describing, among others, multiple cases of forgetting).

<sup>124</sup> The approach is well illustrated by the Lord Chancellor’s speech in *R v. Adomako*:

The jury must go on to consider whether that breach of duty should be characterised as gross negligence and therefore as a crime. This will depend on the seriousness of the breach of duty committed by the defendant in all the circumstances in which the defendant was placed when it occurred. The jury will have to consider whether the extent to which the defendant’s conduct departed from the proper standard of care incumbent upon him . . . was such that it should be judged criminal.

*R v. Adomako* [1995] 1 AC 171 (HL) 187 (Lord Mackay of Clashfern LC) (appeal taken from Eng.).

<sup>125</sup> Moore and Hurd, *supra* note 122 at 148.

<sup>126</sup> Repeated in Judge Magruder: *Ave Atque Vale*, 28 HARV. L. REC. 6, 7 (1959).

of criminal negligence should apply to ordinary negligence, just less. But in many ways this strategy for distinguishing criminal negligence is unsatisfying. Punishment does not at all seem like the same thing as compensatory damages or moral blame, only more; locking someone in prison seems different in kind from ordering him to pay compensation or merely blaming him interpersonally, not just different in degree.<sup>127</sup> If punishment differs in kind so significantly from other responses to negligence, one might expect the forms of negligence that deserve punishment to differ in kind too, not merely in degree, from the forms of negligence that do not. Of course, how grossly a defendant deviates from the standard of care could be one factor in determining his liability: perhaps sufficiently de minimis instances of even culpable negligence should not be punished. But it might seem that the degree of deviation should not be the only distinguishing factor; no arbitrary point in the scale of deviations from the standard of care should determine whether a defendant is subject to punishment or merely to other forms of moral or legal liability.

My account, by contrast, posits a different kind of relationship between criminal and ordinary negligence. Negligence, in general, is commonly understood as the omission of a precaution: as Ken Simons explains, “the primary fault underlying a negligence claim is the actor’s failure to take a reasonable precaution against the risk of harm.”<sup>128</sup> Criminal negligence would not share this feature with ordinary negligence were it a novel form of culpability for inadvertent risk-creation, since it would require only whatever psychological attitude sufficed for that novel form of culpability, whether or not a precaution was omitted.<sup>129</sup> By contrast, in this respect criminal negligence as I analyze it resembles ordinary negligence, for it does require the omission of a certain kind of precaution—namely, the omission of adequate inquiry. It is distinguished from ordinary negligence, in turn, simply because that precaution is omitted with whatever sort of subjective awareness is generally required for criminal culpability. Negligence always involves a breach of the standard of care,

---

<sup>127</sup> Many philosophical definitions of blame remove it exceptionally far from imprisonment: On T.M. Scanlon’s view, for example, “to blame a person for an action . . . is to take that action to indicate something about the person that impairs one’s relationship with him or her, and to understand that relationship in a way that reflects this impairment.” THOMAS SCANLON, *MORAL DIMENSIONS: PERMISSIBILITY, MEANING, BLAME* 122–23 (2008).

<sup>128</sup> Simons, *supra* note 18, at 286.

<sup>129</sup> For discussion, see *id.* at 292–93 & n.25.

but in tort such breaches are not ordinarily evaluated based on the defendant's actual mental states.<sup>130</sup> In evaluating criminal negligence, a jury must instead determine whether the defendant grossly deviated from the standard of care "considering the nature and purpose of his conduct and the circumstances known to him."<sup>131</sup> In a literal sense, then, my account analyzes criminal negligence as negligence that is criminal: it is negligent for the same reason as any other instance of negligence—because a precaution was omitted—and it is criminal for the same reason as any other crime—because of the defendant's culpable awareness.<sup>132</sup> This analysis strikes me as identifying a more compelling and less arbitrary basis for distinguishing criminal and ordinary negligence, which better fits the very non-arbitrary distinction between criminal punishment and other remedies: the psychological condition that justifies criminal punishment for negligence is the very same condition that justifies criminal punishment in general.

Though my account does in general analyze criminal negligence simply as negligence that is criminal, it does impose one restriction on the scope of criminal negligence, in addition to the requirement of culpability, that does not limit the scope of ordinary negligence. While the omission of any kind of precaution can be negligent, I have interpreted negligence to be criminal only when it involves the omission of a very particular kind of precaution—inquiry. If criminal negligence is simply negligence that is criminally culpable, why would it be limited to the culpable omission of this precaution, rather than extending to the culpable omission of any precaution? Some of my arguments would be consistent with extending liability to other omitted precautions, and some such extensions may be justified. But there are also some good reasons for limiting criminal liability to this particular kind of

---

<sup>130</sup> Tort law requires a defendant to recognize a risk "if a reasonable man would do so while exercising (a) such attention, perception of the circumstances, memory, knowledge of other pertinent matters, intelligence, and judgment as a reasonable man would have," regardless of whether the defendant's actual attention, perception, memory, knowledge, intelligence, or judgment fell short. *See* RESTATEMENT (SECOND) OF TORTS § 289 (1965).

<sup>131</sup> MODEL PENAL CODE § 2.02(2)(d) (AM. LAW INST. 1962).

<sup>132</sup> Indeed, while on one traditional approach the distinction between criminal and civil negligence does turn on whether the negligence was "gross," another distinguishes them along the lines I defend here. As the commentaries to the Model Penal Code explain, earlier penal codes had often defined negligence simply as ordinary negligence that is criminally culpable. MODEL PENAL CODE § 2.02 cmt. 4. New York, for example, defined negligence as "a want of attention to the nature or probable consequences of the act or omission as a prudent man ordinarily bestows in acting in his own concerns." NEW YORK PENAL LAW § 3(1) (Albany, J.B. Lyon Company 1909). Many particular offenses were then defined to require "culpable negligence." *E.g.*, *id.* § 1052(3) (manslaughter).

omission. Criminal law generally hesitates to punish omissions: under the Model Penal Code, for example, they suffice for liability only if a duty to avoid a particular omission is explicitly imposed by law, criminal or otherwise.<sup>133</sup> Criminal liability for the omission of inquiry, then, is itself exceptional. And liability for omitting inquiry is justified by some reasons that do not generalize to other precautions.

Inquiry is a unique precaution. Most precautions apply only to a small subset of our actions, but inquiry is ubiquitous: we must possess some information about every decision we make in order to reason at all, and thus the possibility of inquiry is relevant to every decision (though of course the amount of information we should acquire can vary greatly). For two reasons, this fact justifies imposing criminal liability for the omission of inquiry. First, because inquiry is ubiquitous, its omission threatens greater harm in total: simply because inquiry is required more often than other precautions, more opportunities exist for it to be omitted, and thus prohibiting it would avoid more harm to aggregate social interests. The greater harm this prohibition prevents constitutes a stronger reason to impose it. Second, because inquiry is ubiquitous, a primary obstacle to prohibiting omissions generally does not apply in the case of inquiry. Everyone constantly omits an enormous range of actions: there are always countless things that an individual could but does not do. Because so many actions are available—many more than could possibly be considered—it often seems unreasonable to expect individuals to realize that one particular action is required on a particular occasion, whereas they can more easily identify whether the single act they perform is prohibited.<sup>134</sup> The law addresses this problem by providing notice if omissions are prohibited: a legal duty requiring an action, for example, provides notice of the required act.<sup>135</sup> But no device is needed to provide notice that inquiry may be required. If a precaution were irrelevant to most decisions, individuals might reasonably fail to realize its relevance on a rare occasion when it was

---

<sup>133</sup> MODEL PENAL CODE § 2.01(3).

<sup>134</sup> As Graham Hughes puts it, “In an offense of commission, the mind of the actor is almost always to some extent addressed to the prohibited conduct . . . With omissions, the great difficulty is that the mind of the offender may not be addressed at all to the enjoined conduct . . . .” Graham Hughes, *Criminal Omissions*, 67 YALE L.J. 590, 600–01 (1958).

<sup>135</sup> See MODEL PENAL CODE § 2.01 cmt. 3 (explaining that the requirement “best deals with most situations in which criminal liability would clearly be appropriate and also gives adequate warning of what omissions may be the basis of such liability.”).

required.<sup>136</sup> It is never reasonable not to realize that information is relevant to choice, however, because information is always relevant to choice. Individuals do not need to be put on notice that inquiry may be required before a particular action, because they should always be on notice that inquiry may be required before every action. Thus, while the importance of notice does caution against extending culpability for negligence to all omitted precautions, it does not similarly caution against imposing liability for the failure to inquire.

#### D. Tracing

In some respects, the analysis I have provided of criminal negligence resembles an existing analysis of negligence—tracing.<sup>137</sup> Like my account, tracing accounts explain culpability for inadvertent risk-taking by arguing that individuals may in some way be culpable for their inadvertence. Furthermore, they too analyze culpability for inadvertence in terms of the risk created by failing to take certain precautions before acting. It might be unclear, then, how my proposal differs from tracing. Nonetheless, one key difference exists: tracing analyzes negligent conduct as requiring defendants to make two distinct decisions to perform two distinct actions.<sup>138</sup> One, of course, is the negligent act itself, a risky action performed without awareness of that risk.<sup>139</sup> Second, though, tracing accounts posit a temporally prior act in which the agent failed to learn relevant information or otherwise take a precaution against his subsequent performance of the negligent act.<sup>140</sup> For example, he might have failed

---

<sup>136</sup> In addition to doctrines concerning liability for omissions, mens rea doctrines themselves might exculpate defendants who fail to realize that an action is required in certain circumstances. If criminal negligence constituted the culpable omission of any precaution, not just inquiry, then a defendant who failed to realize the relevance of a precaution to a particular decision still might avoid liability on a narrow conception of subjective awareness. But other notions of subjective awareness require only tacit or dispositional beliefs. And it might seem unreasonable to expect an individual merely tacitly or dispositionally aware of the risk created by omitting a precaution to recognize its importance in a particular context absent some notice that the precaution was relevant in that context.

<sup>137</sup> For comprehensive articulations of the tracing accounts, see Holly Smith, *Culpable Ignorance*, 92 PHIL. REV. 543 (1983); Michael J. Zimmerman, *Negligence and Moral Responsibility*, 20 NOÛS 199 (1986).

<sup>138</sup> Smith, *supra* note 137, at 547 (“[T]he relevant cases all involve a sequence of acts: an initial act, in which the agent fails to improve [or positively impairs] his cognitive position, and a subsequent act in which he does a wrong because of his resulting ignorance.”); Zimmerman, *supra* note 137, at 200 (“Bert must [i] not have adverted to the possibility of causing damage or injury at the time he threw the brick, but [ii] have adverted to this possibility earlier.”).

<sup>139</sup> Smith, *supra* note 137, at 545 (“[W]e must make sure that the wrongful act is justified relative to the agent’s actual beliefs at the time he performs it [even though it is not justified relative to the beliefs he ought to have had].”); Zimmerman, *supra* note 137, at 200 (“[A]t the time he threw the brick, Bert did not advert to . . . the possibility that he might thereby cause damage or injury.”).

<sup>140</sup> Smith, *supra* note 137, at 547 (describing “an initial act . . . in which the agent fails to improve [or positively impairs] his cognitive position”); Zimmerman, *supra* note 137, at 201–02 (“P’s action is neglectful only if he

to perform a particular bit of research that later proves relevant,<sup>141</sup> or he might have failed to erect a physical barrier to protect others from his thoughtless behavior.<sup>142</sup> The failure to take that precaution, furthermore, must have been culpable.<sup>143</sup> If all these conditions are met, tracing accounts conclude, the agent is culpable not just for omitting the precaution but also for later inadvertently causing harm: the earlier culpable omission explains why the agent is culpable for the subsequent negligent act.<sup>144</sup>

Many of the criticisms levied against tracing depend on its division of negligence into two distinct acts. First, tracing accounts rely on a derivative notion of culpability: the subsequent act is culpable because of the culpability of the prior omission. But even when some prior culpable omission is causally responsible for the subsequent negligent act, many have been skeptical that the kind of culpability involved in the prior act will generally be of the right sort to justify holding a defendant culpable for the latter act.<sup>145</sup> In general, performing a culpable wrong does not make an individual responsible for all of the results it happens to cause.<sup>146</sup> Culpable risk-creation ordinarily makes individuals responsible only for consequences that in some sense fall within the scope of the risk that was culpably

---

decides not to take both [1] a precautionary measure and [2] any probably more effective precautionary measure.”).

<sup>141</sup> *E.g.*, Smith, *supra* note 137, at 550 (describing a case in which an individual fails to conduct a blood-type survey for a science project and subsequently, after an accident, cannot tell paramedics the victim’s blood type).

<sup>142</sup> Zimmerman, *supra* note 137, at 202 (giving, as an example, a bricklayer’s constructing a protective screen to block any defective bricks he might, from habit, discard over his shoulder).

<sup>143</sup> Smith, *supra* note 137, at 548 (“[The prior] act must be more than objectively wrong: it must also be one for which the agent is culpable.”); Zimmerman, *supra* note 137, at 202–03 (requiring the agent to unjustifiably fail to take a precaution).

<sup>144</sup> Smith, *supra* note 137, at 570 (“In cases of culpable ignorance, the unwitting act is a risked upshot of the benighting act, so the agent is to blame for it . . . .”); Zimmerman, *supra* note 137, at 206 (“[T]he inadvertence at the time . . . of the omission renders the omission beyond the agent’s control at that time, but this fact does not render the omission beyond the agent’s control entirely. On the contrary, the control requisite for moral responsibility may be, and often is, anchored in the earlier advertence . . . .”).

<sup>145</sup> In addition, many have doubted whether the prior culpable acts required for the tracing strategy to succeed actually exist in most cases of negligence. *See* Matt King, *The Problem with Negligence*, 35 SOCIAL THEORY AND PRACTICE 577, 580 (2009) (“Tracing plainly will not work, however, in cases of negligence, for in such cases it is difficult to demonstrate what the initial choice is.”); Manuel Vargas, *The Trouble with Tracing*, 29 MIDWEST STUD. PHIL. 269, 275 (2005) (describing cases in which the “source of behavior was acquired or retained under conditions where the agent could not have reasonably foreseen the later consequences of having that disposition, habit, or character trait”). Of course, advocates of tracing could concede that these cases of negligence are not culpable.

<sup>146</sup> Finkelstein, *supra* note 4, at 592 (“And I would exclude all unforeseen effects from the scope of responsible agency.”); Vargas, *supra* note 145, at 274 (proposing a foreseeability requirement to determine the scope of responsibility for consequences).



created.<sup>147</sup> And in most cases it seems doubtful that the risk created by an initial, culpable act will ordinarily encompass the particular risk of causing harm through the subsequent, inadvertent act actually performed—someone who breaks a promise to his mother to read a particular book that by chance contains facts relevant to a later decision is hardly culpable for making that decision incorrectly.<sup>148</sup> Even if the subsequent act does fall within the scope of the risk, furthermore, tracing faces a second objection. It argues that an agent’s culpability for a prior omission can make him culpable for a subsequent act. But the nature of this transferred culpability seems quite mysterious—just how, exactly, does culpability for one act somehow carry over to a subsequent act that, by stipulation, lacks culpability in its own right?<sup>149</sup> Perhaps, some opponents of tracing concede, an earlier culpable act may be reckless in its own right, but that culpability is limited to the prior act; it cannot provide a novel explanation for the culpability of the distinct, subsequent act.<sup>150</sup>

Both of these objections depend on the tracing account’s insistence on explaining the structure of negligence in terms of two distinct decisions to perform two distinct acts, such that the culpability of the former can ground the culpability of the latter. Only because there are two distinct acts can the question arise of whether the harm caused by the subsequent act falls within the risk created by the prior act; only because there are two distinct acts can the question arise of whether the culpability involved in the prior act can somehow be transferred to the latter act. One obvious way to avoid these objections, then, is simply to dispense with the requirement of two acts. My account of criminal negligence adopts precisely this approach: on my view, the negligent agent faces a single choice between acting, inquiring, and not acting, and he is culpable for how he makes that decision.<sup>151</sup> Of

---

<sup>147</sup> See MODEL PENAL CODE § 2.03(3) (AM. LAW INST. 1962); Rosen, *supra* note 12, at 603–04; Smith, *supra* note 137, at 550–51.

<sup>148</sup> Rosen, *supra* note 12, at 603–04; see also Moore & Hurd, *supra* note 122, at 179–80; Vargas, *supra* note 145, at 274–82.

<sup>149</sup> E.g., Smith, *supra* note 137, at 559 (“But the fact that [the agent] earlier had faulty motives does not show that he now has faulty motives.”). For a contrary argument defending such transfers, see Rosen, *supra* note 12, at 601–03.

<sup>150</sup> On this view, for example, the defendant in *People v. Decina*, 138 N.E.2d 799 (N.Y. 1956) was reckless, not negligent, when, knowing that he might have an epileptic seizure that could cause him to crash, he nonetheless drove a car and killed multiple pedestrians. ALEXANDER, FERZAN, & MORSE, *supra* note 4, at 80–81. Of course, even if only the prior act is culpable, a further question may remain concerning how culpable it is—particularly whether the occurrence of subsequent harm increases its culpability. See Smith, *supra* note 137, at 562–68.

<sup>151</sup> In fact, Holly Smith does discuss this sort of case, too:

course, time will elapse between his decision and the harm that it causes, if it causes harm, just as time usually elapses between a culpable decision and its harmful results. But though subsequent events may affect whether an action satisfies the elements of an offense, his culpability depends on his failure to inquire alone.

Because my account does not posit two distinct choices, it avoids the objections raised against traditional versions of tracing. There is no need to explain how culpability somehow transfers from a prior culpable act to a subsequent inadvertent one if negligence involves only a single act. And because of how I have defined the culpable choice involved in negligence, I can explain why the harm caused does fall within the scope of the risk: when inquiry is required because it would reduce the risk that an agent will inadvertently perform an excessively risky action, then a harm clearly does fall within the scope of the risk if it occurs only because the agent's failure to inquire prevented him from recognizing the risk his action created. My account, then, vindicates the key insight of tracing—namely, that the failure to inquire can sometimes itself constitute a source of culpability—while avoiding the objections that arise from the particular logical structure that advocates of tracing have employed to analyze that failure.

#### **IV. Conclusion**

The standard understanding of criminal negligence is paradoxical: it holds individuals liable for inadvertently creating risks, even though ordinarily advertence of some sort is a necessary requirement for liability to be justified. In this chapter, I have attempted to dissolve this paradox rather than resolving it: I do not pick sides in the conflict between the requirement of awareness and the intuition that particular acts of negligence are culpable but rather show that they need not conflict. My

---

[T]he driver . . . should have checked his mirror earlier, but given that he did not, he should check it now rather than back down the driveway. There are many cases in which enquiry should be made earlier, but it is better to enquire now rather than act without its benefit. In such a case the agent is indeed culpable for his act, but the culpability might be wholly traceable to the fact that he knows himself to be performing an act less good than its alternative, namely conducting further enquiry.

Smith, *supra* note 137, at 546. Nonetheless, she mentions it only to set it aside before presenting and evaluating the standard version of tracing. By contrast, on my view this case is a paradigmatic example of criminal negligence.

account reconceptualizes criminal negligence by understanding it to regulate a distinctive aspect of conduct—namely, the inquiry individuals perform before acting. The standard approach understands criminal negligence as a modification of recklessness: it is the creation of a substantial and unjustifiable risk, minus the awareness of risk that makes recklessness culpable. By contrast, I understand it as the criminal analogue of ordinary negligence: it is the omission of a key precaution, inquiry, plus the awareness of risk that generally justifies the imposition of criminal liability. Thus, criminal negligence plays a distinctive role, complementary to recklessness, in regulating the risks that actions create: it requires individuals to take steps to reduce risk in circumstances when actions are not dangerous enough to justify prohibiting them entirely. And because negligence regulates the risk created by the failure to inquire, it can involve both culpable advertence towards that risk and inadvertence towards the risk that would have been revealed by the omitted inquiry.