

Torture and Moral Integrity

A Philosophical Enquiry

By

MATTHEW H. KRAMER

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© M Kramer 2014

The moral rights of the author have been asserted

First Edition published in 2014

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Crown copyright material is reproduced under Class Licence
Number C01P0000148 with the permission of CPSI
and the Queen's Printer for Scotland

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10017, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2013957961

ISBN 978-0-19-871420-0

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

1

Introduction I: Moral Conflicts and Deontology

We are oft to blame in this, 'tis too much proved, that with devotion's visage
and pious action we do sugar o'er the devil himself.

[*Hamlet*, III.i.46–9.]

This book is about the wrongness of torture and about the nature of morality. It deals at length with multiple types of torture, and it seeks to explain why the types that figure most prominently in contemporary philosophical discussions are always morally wrong. At the same time, it plumbs the general structure of morality and the intricacies of moral conflicts, and it probes some of the chief grounds for the moral illegitimacy of various modes of conduct. It tackles a concrete moral problem—a problem heatedly debated during recent decades in the governmental and military institutions of many countries as well as in academic circles—and it likewise tackles some very abstract issues in moral and political philosophy. Moreover, as will become apparent at numerous junctures, the abstract ruminations and the concrete prescriptions are closely connected. My abstract reflections on the structure of morality are the vital background for my approach to torture, and my approach to torture is a natural outgrowth of those abstract reflections.

In keeping with the overall character of this book, its Introduction is divided into two chapters. This opening chapter will address some important matters in the more abstract reaches of moral philosophy—as it disambiguates several key concepts in order to clarify the import of moral conflicts, and as it elucidates the distinction between deontological obligations and consequentialist obligations. Along the way, it will delineate the general structure of morality. Chapter 2 will then move to a more concrete level of philosophizing, by proceeding to recount many of the complications that surround any attempt to distil the defining characteristics of torture. These discussions in the first two chapters, at quite differing tiers of abstraction and concreteness, will form the basis for the subsequent chapters' treatment of the wrongness of torture and the value of moral integrity. Chapter 3 will endeavour to explain why torture is morally wrong; Chapter 4 will defend my absolutist stance on torture against a powerful consequentialist critique; and Chapter 5 will prescribe how the law in any jurisdiction should address the matter of torture. Both in those later chapters and in these introductory chapters,

the mutual supportiveness of the abstract argumentation and the concrete argumentation in this book will become evident.

1.1. Moral Conflicts

To come to grips philosophically and morally with the problem of torture, this book needs to examine the phenomenon of moral conflicts. Throughout the volume, I usually take a moral conflict to be a situation in which some person P is under a moral duty-to- ϕ and simultaneously under a moral duty-not-to- ϕ .¹ However, in line with most other philosophers who write on these issues, I also apply the phrase ‘moral conflict’ to any situation in which P simultaneously bears moral duties whose contents are contraries rather than contradictories. A situation of the latter sort obtains when P is under a duty to do x and simultaneously under a duty to do y , where (1) ‘ P does x ’ entails the negation of ‘ P does y ’ and (2) ‘ P does y ’ entails the negation of ‘ P does x ’ and (3) the negation of ‘ P does x ’ is logically consistent with the negation of ‘ P does y ’. Though the clashing duties within a moral conflict can perfectly well coexist, they can never be jointly fulfilled: the fulfilment of either of them entails the non-fulfilment of the other.

1.1.1. The disambiguation of some key concepts

Before we directly ponder a number of queries that have been raised about the occurrence of moral conflicts, this chapter should disambiguate some key concepts that figure saliently in discussions of these matters and specifically in discussions of torture. An examination of several overlapping and oft-observed distinctions will not merely help to avert confusion, but will likewise illuminate a number of the points of contention that surface in the later portions of this book. Attentiveness to those distinctions will greatly enhance one’s understanding of the nature of moral conflicts.

1.1.1.1. Two senses of ‘prima facie’

As I have observed elsewhere (Kramer 1999a, 267–9; 2004, 290–1; 2011, 53–4), the phrase ‘prima facie’ is construable in two quite divergent ways. On the one hand, it can mean ‘presumptively’ or ‘upon initial examination’. When the phrase is used in this epistemic sense, a person who declares that some state of affairs prima facie obtains is saying that an initial inspection of the matter indicates that the specified state of affairs does obtain. As Susan Hurley wrote: ‘*Prima facie* reasons are like rules of thumb, that give us reasons provisionally but may turn out not to apply when we learn more about the situation at hand, in which case they have no residual reason-giving force’ (Hurley 1989, 133). This pattern of usage

¹ The variable ‘ ϕ ’ can stand for any verb(s) or verb phrase(s), denoting any action(s) or omission(s).

carries over into the workings of legal institutions, where conclusions that have prima facie been demonstrated are conclusions that are tentatively deemed to be demonstrated in light of the incomplete evidence that has been submitted. (In the law, the notion of prima facie evidence is connected with issues relating to the distribution of the burden of proof between plaintiffs and defendants.) Insofar as this pattern of usage is regnant in moral and political philosophy, theorists who discuss prima facie moral duties are referring to properties or states that are believed upon first inspection to be moral duties. No such property or state will qualify as a full-fledged moral duty unless the initial consideration of the factors that bear on its existence as a duty is confirmed by subsequent investigations of all further relevant factors (if any). In short, when things are classified as prima facie moral obligations in this fashion, only some of them are veritable moral obligations. Many are mere appearances that turn out not genuinely to be moral obligations when all things have been considered.

On the other hand, 'prima facie' can mean 'pro tanto' or 'susceptible to being overtopped in importance by some competing factor'. When the phrase is used solely in this second way, a prima facie moral duty is a genuine moral duty that might be exceeded in stringency or importance by a countervailing moral imperative. In this vein, a *merely* prima facie obligation is contrasted with an overtoppingly stringent obligation. That is, the relevant dichotomy lies between any moral duty that does not exceed every competing moral requirement in importance and any moral duty that does. Here the prima facie status of a moral duty has nothing to do with first appearances or tentative identifications. If someone is under a prima facie moral obligation in this sense and not in the epistemic sense, he is under it as a result of all relevant considerations. Inasmuch as 'all things considered' is construed as an epistemic phrase that means 'in accordance with what would be found by a thorough investigation', every moral duty that is prima facie solely in the sense of not being overtoppingly stringent in some possible situations (rather than in the sense of being an initial appearance) is an all-things-considered moral duty. The standing and scope of any such duty as a moral obligation are determined by the full array of circumstances in which the obligation obtains.

This point has sometimes been obfuscated, not least by David Ross's suggestion that any prima facie moral obligation consists in a '[t]endency to be one's duty' (Ross 1930, 28). Such a formulation promotes confusion, for it fails to distinguish between a tendency to be a duty of any sort and a tendency to be an overtoppingly stringent duty.² Any number of considerations can engender a tendency of the former kind. For example, the pain and indignity that would be suffered by Edward in the event of his undergoing a thrashing by Frank are factors that militate in favour of Frank's being morally obligated to forbear from thrashing Edward in any particular context. Those factors do not always produce or sustain such an

² This distinction is missed by Joseph Raz when he draws a parallel between (i) the liberty of courts in some circumstances to overrule precedents which they would in other circumstances be obligated to apply and (ii) the overtoppingly stringent duty of a promisor to decline to abide by her promissory obligation in some circumstances. See Raz 1979, 114.

obligation, however. If Edward himself launches a serious and unprovoked assault against Frank, then Frank is morally at liberty to use as much force as is necessary to fend off Edward's attack. He is not under any moral obligation at all—not even a merely prima facie duty, much less an overtoppingly stringent duty—to abstain from landing blows on Edward in such circumstances. At first glance we might think that Frank is under at least a merely prima facie moral duty to restrain himself from using violence against Edward, but a more detailed exploration of his predicament reveals that not even a merely prima facie moral duty is applicable. There are factors that tend toward the existence of a duty, but none of them actually eventuates in it. Such a situation differs markedly, then, from a situation in which a veritable prima facie moral duty does exist while being overtopped by a weightier moral requirement. A veritable prima facie moral obligation does not *tend* toward the existence of an obligation; it *is* a moral obligation. What it tends toward is the existence of an overtoppingly stringent moral obligation. If that latter tendency is fulfilled in any particular context, the moral duty in question prevails over all conflicting moral requirements. If the tendency toward overtopping stringency is instead checked by the existence of some competing moral requirement that is more important, the prima facie duty still obtains as such. Even when overtopped in importance, it is indeed a moral obligation rather than a bare appearance or a sheer factor (like Edward's pain) that contributes toward the existence of such an obligation.

Henceforth, whenever this book uses 'prima facie' without any explicit signal that those words are to be construed in the first sense above, it is employing the phrase in the second sense. That is, instead of being perceived as an epistemic caveat, the phrase 'prima facie' (in the absence of any explicit signal of the sort just mentioned) should be understood herein as an indication that a duty under discussion is susceptible to being surpassed in normative importance by some conflicting moral requirement. When a moral obligation is indeed so surpassed, it is *merely* prima facie—which, again, does not detract at all from its status as a fully genuine moral obligation that cannot be breached without the incurring of remedial duties.

1.1.1.2. Weak permissibility versus strong permissibility

In application to human conduct, permissibility and wrongness are contradictories. Any type or instance of human conduct is permissible if and only if it is not wrong. Now, if some person *X* is permitted to perform some action *q*, then *X* is not obligated to not perform *q*. (To avoid any ambiguities in my prose, my placement of 'not' in several of the sentences in this subsection will create some ugly split infinitives.) So much is clear, but we need here to take account of two ways in which someone can be morally obligated to do or not do something.

1.1.1.2.1. Two types of obligations and two types of permissibility

The germane distinction, which has already been invoked at a few junctures above, is between overtoppingly stringent and non-overtoppingly stringent moral

obligations.³ An overtoppingly stringent moral requirement exceeds in importance all the moral duties that run counter to it, or is unopposed by any competing moral duties. A non-overtoppingly stringent moral requirement *R* does not exceed in importance all the moral duties that run counter to it. (Any competing moral requirements might be equal in importance to *R*, or they might exceed it in importance, or they might be insusceptible to any determinate comparisons with it because of problems of incommensurability.)

With reference to these two broad types of obligations, we can apprehend two broad types of permissibility:

Weak Permissibility. Some person *X* is permitted to perform some action *q* if and only if *X* is not under any overtoppingly stringent obligation to not perform *q*.

X's being weakly permitted to perform *q* is consistent with the proposition that *X* is under a non-overtoppingly stringent obligation to not perform *q*. In that respect, weak permissibility differs from strong permissibility.

Strong Permissibility. *X* is permitted to perform *q* if and only if *X* is neither under an overtoppingly stringent obligation to not perform *q* nor under a non-overtoppingly stringent obligation to not perform *q*.

X's being strongly permitted to perform *q* is inconsistent with the proposition that *X* is under a non-overtoppingly stringent obligation to not perform *q*. Note that the proposition 'X is strongly permitted to perform *q*' entails the proposition 'X is weakly permitted to perform *q*', but not vice versa.⁴

Whenever this book uses the term 'permissible' or 'permissibility' without any qualification, it is invoking the notion of permissibility in the strong sense. Consequently, the term 'impermissible' is to be understood herein as 'not strongly permissible'. That is, I generally assume that an action *q* is impermissible for a person *X* unless *X* is not under any obligation whatsoever to not do *q*. For the purpose of gauging whether any type or instance of conduct is impermissible, this book will generally not discriminate between situations in which someone is under an overtoppingly stringent obligation to eschew *q* and situations in which someone is only under a non-overtoppingly stringent obligation to eschew *q*. Precisely because my ascriptions of impermissibility will generally not discriminate between those two kinds of situations, this book will contend that morally optimal courses of action can be morally impermissible. If two moral duties clash, and if one exceeds

³ See Kramer 2004, 280–1. Some closely related distinctions are sustainedly brought to the fore in Sinnott-Armstrong 1988. Although the term 'overriding' is much more common than 'overtopping' in discussions of these matters, I disfavour the former term because it conveys the impression that a less important duty is eliminated or cancelled in any conflict with a more important duty. We shall explore this point shortly.

⁴ Michael Moore has also posed a distinction between weak permissions and strong permissions (Moore 2007, 42–4), but his distinction is markedly different from mine. Likewise very different from my own distinction is the contrast between weak permissions and strong permissions that was drawn half a century ago by G. H. von Wright (1963, 86). In the text here and elsewhere, I am invoking only my own weak/strong dichotomy. Although I have expounded that dichotomy here with reference to actions, it is equally applicable—*mutatis mutandis*—to omissions.

the other in importance, then compliance with the former is a morally optimal but impermissible course of conduct. The moral optimality of such compliance will extenuate but not eliminate the breach of duty involved; it will therefore not eliminate the impermissibility of the compliant course of conduct. These attributions of impermissibility are sustainable because I generally rely on the strong conception of permissibility. (Throughout this book, I use 'legitimate' interchangeably with 'permissible' and 'illegitimate' interchangeably with 'impermissible'.)

1.1.1.2.2. Infringements versus violations

Note that this exposition of the two types of permissibility can refine one's understanding of a distinction which was originally expounded by Judith Jarvis Thomson (and slightly later by Joel Feinberg) and which has been quite conspicuous in political and moral philosophy ever since: the distinction between infringements and violations of rights.⁵ Under the terms of that dichotomy, all violations are infringements, but not all infringements are violations. An infringement occurs when someone has a right that p be the case—where ' p ' designates a certain proposition—and someone else causes p to be false. For example, if Henry has a right that Ron not bang him in the nose, then Ron infringes the right if he bangs Henry in the nose. A violation is a culpable infringement. If Ron's banging of Henry's nose is attributable to malice or recklessness or carelessness, then the infringement of Henry's right is a violation thereof. Contrariwise, if the banging occurs despite Ron's careful efforts to avoid it, the infringement is a mere infringement.

Proponents of the infringement/violation distinction generally maintain that mere infringements are morally permissible. Only violations are morally impermissible. In other words, these philosophers take culpability to be a necessary condition for impermissibility and thus for wrongdoing. Now, given that an infringement is a contravention of a moral right and is thus a breach of a moral duty, these philosophers may seem clearly mistaken in contending that infringements are morally permissible. Were they adhering to this book's terminological patterns, they would indeed be mistaken. However, as becomes evident when one peruses their relevant writings, the exponents of the infringement/violation duality have typically relied (at least implicitly) on the weak conception of permissibility when submitting that mere infringements of people's rights are permissible. They maintain that, so long as any person X is conducting himself in accordance with all the overtoppingly stringent moral duties that are incumbent upon him, X is conducting himself permissibly. The fact that he might be acting athwart a non-overtoppingly stringent moral duty—for example, by breaking an engagement for lunch in order to go to the aid of a seriously injured pedestrian—is compatible with the weak permissibility of his course of action. In their repeated affirmations of the permissibility of mere infringements, then,

⁵ See Thomson 1986, chaps 3–5; Thomson 1990, 122; Feinberg 1980, 229–32. For some discussions of the infringement/violation dichotomy, see Botterell 2008; Cane 2002, 107; Coleman 1992, 282–3, 299–302; Fletcher 1985, 977; Fletcher 1993, 175, 177; Lee 2012, 144–6; McConnell 1996, 42; Oberdiek 2004; Oberdiek 2008; Parent 1980, 406–8; Rainbolt 2006, chap. 6; Simester 2008, 297–8; Sinnott-Armstrong 1988, 51–2.

Thomson and Feinberg and their followers are not really committing any errors. Instead, they are exhibiting their adherence to a conception of permissibility that diverges from my own conception.

Still, although the defenders of the infringement/violation contrast are not committing outright errors, their understanding of permissibility as weak permissibility should not commend itself to anyone who wants to grasp the complexities of the moral bearings of conduct such as torture in extreme situations. After all, irrespective of whether a moral duty is overtoppingly stringent or non-overtoppingly stringent, any breach of it will give rise to remedial obligations. In the scenario mentioned just above, for instance, the rescuer of the injured pedestrian will have incurred a moral obligation to remedy his breach of a promissory duty (most likely through an apology and a brief explanation). Of course, to say as much is scarcely to suggest that the division between overtoppingly stringent obligations and non-overtoppingly stringent obligations will play no role whatsoever in my analyses. For example, that division plays a key role in determining the *sizeableness* of the remedy that is required in the aftermath of a breach of some moral duty. A breach undertaken in order to avoid a transgression of an overtoppingly stringent moral duty will be extenuated more heavily than a breach undertaken in order to avoid a transgression of a non-overtoppingly stringent moral duty. Nevertheless, the need for a remedy of *some* sort in the aftermath of the contravention of a moral duty is unaffected by the status of the contravened duty as overtoppingly stringent or non-overtoppingly stringent. Because of the irrelevance of the overtopping/non-overtopping distinction in that crucial respect, this book adheres to the strong conception of permissibility throughout (except when the weak conception is clearly signalled, usually as I am expounding someone else's views).

1.1.1.3. *Strong justification versus weak justification*

Closely cognate to the distinction between strong permissibility and weak permissibility is another dichotomy, between strong justification and weak justification (Kramer 2011, 54). Suppose that some person P asserts that the use of torture as a technique of interrogation can be morally justified in certain situations. On the one hand, P might be declaring that the use of interrogational torture can sometimes be strongly permissible as well as morally obligatory. That is, P might be contending that such a technique in some possible circumstances will not be in violation of any moral obligations. On the other hand, P might instead be declaring that moral obligations which sometimes require the use of torture as an interrogational technique are so dauntingly weighty in certain contexts that they can exceed in ethical importance the potentially stringent moral obligations which forbid any such use. Construed in this latter way, P's assertion would not be denying that the use of torture for interrogation is always morally impermissible; P would simply be maintaining that that morally impermissible mode of conduct can in some imaginable circumstances be morally imperative.

Understood in the first way just recounted, P's assertion about torture is a strong justification. Understood in the second way, it is a weak justification. In other

words, if a mode of conduct is strongly justified, it is both morally obligatory and morally permissible—strongly permissible—in the circumstances in which it occurs. If a mode of conduct is only weakly justified, it is morally obligatory and morally optimal but not morally permissible (that is, not strongly permissible). Whereas both any strongly justified course of conduct and any weakly justified course of conduct are morally optimal,⁶ only the former is morally permissible. A weakly justified course of conduct is morally wrong, even though its wrongness is not graver than that of any opposing course of conduct in the circumstances.⁷

1.1.1.4. *Two senses of 'rightness'*

As a moral term, the adjective 'right' can be either an antonym of 'wrong' or a synonym of 'correct'. Although those two senses of the adjective are not mutually exclusive, they are not equivalent. If a course of conduct is right in the first sense, it is strongly permissible; engaging in such a course of conduct will not breach any moral duties and will therefore not amount to a wrong at all. If some course of conduct CC is right in the second sense, then it is morally optimal and is thus at least weakly permissible, but it might not be strongly permissible. Suppose that a person is under a duty to engage in CC and simultaneously under a duty not to engage in it, and suppose that CC is the *uniquely* optimal course of conduct for the person in the circumstances. Albeit the former obligation is more stringent than the latter, the latter duty continues to exist as such. Accordingly, although CC is right in the sense of being the morally correct thing to do, it is not right in the sense of being strongly permissible. Because the person's adoption of such a course of conduct involves the contravention of his moral duty to refrain from adopting it, CC is morally wrong. It is uniquely optimal—the non-adoption of CC would be an even graver wrong—but it is only weakly permissible.

1.1.1.5. *Two senses of 'absolute'*

In debates over the moral bearings of torture, one of the frequent points of controversy is whether the moral prohibition on torture is absolute. That much-broached question is in need of disambiguation, however. In posing that question, one might be asking whether the duty not to engage in torture is always and everywhere

⁶ Though the term 'optimal' (or the term 'best') is evaluative, it is not to be understood along consequentialist lines here or elsewhere in this book. Optimality is instead a function of deontic stringency. If someone's moral duty-to- ϕ conflictingly coexists with her moral duty-not-to- ϕ , and if the stringency of the former duty is greater than that of the latter, her fulfilment of the duty-to- ϕ is the lone morally optimal course of conduct for her in the circumstances. If the two duties are instead equally stringent, then her fulfilment of either of them is morally optimal in the circumstances. If her duty-to- ϕ is not countervailed by any conflicting moral requirements, compliance with that duty is both morally optimal and strongly permissible.

⁷ Note that the division between strong justifications and weak justifications does not encompass all courses of conduct. It obviously excludes morally unjustified courses of conduct, but it also omits all non-obligatory courses of conduct (many of which, such as my scratching of my nose in virtually any ordinary circumstances, are not morally unjustified).

binding in all possible worlds, or one might be asking whether the stringency of that ever-binding duty always and everywhere overtops the stringency of any duties that might conflict with it. Although an affirmative answer to the latter version of the question entails an affirmative answer to the former version, there is no entailment between affirmative answers in the other direction. Strong absoluteness entails weak absoluteness, but not vice versa.

If the moral prohibition on torture is strongly absolute, it is not only binding at all times in all places in all possible worlds; in addition, it is also always of greater normative importance than any possible countervailing moral requirements. One's insistence on the absoluteness of that prohibition in this strong sense is perfectly consistent with one's recognition that there can indeed be countervailing moral requirements that militate against compliance with the prohibition. In other words, a strongly absolute moral obligation can be locked in conflicts with competing moral obligations that are inferior in their normative strength.⁸ Although those latter obligations are indeed less stringent, they impose genuine moral demands that will give rise to remedial duties if they are left unfulfilled. In such circumstances, then, compliance with the strongly absolute moral prohibition is not strongly permissible; it is only weakly permissible.

Any weakly absolute moral prohibition WM is binding everywhere and always in all possible worlds. No circumstances, however exigent, can ever negate or diminish the demands of such a prohibition. Nonetheless, although WM is irrepressibly operative in all possible situations, it can be locked in conflicts with competing moral requirements (just as strongly absolute moral prohibitions can be). Moreover, in some imaginable circumstances, the competing moral requirements are more stringent than WM with which they conflict. In such circumstances, compliance with WM would not be even weakly justified. There is not available any strongly justified course of conduct in such a setting, and the sole weakly justified course of conduct resides in breaching WM. Any such breach is indeed a breach—that is, a moral wrong—and it will thus trigger remedial obligations.

Of course, if WM is formidably stringent, there might never materialize any actual situations in which the contravention of WM would be morally optimal. Nevertheless, so long as there *could* arise some context in which such a contravention would be optimal, WM is only weakly absolute rather than strongly absolute. Even if conformity to WM is always morally optimal in the actual world, there are never any guarantees of that optimality—for there remain counterfactual scenarios in which some actions at odds with WM are the morally best modes of conduct (the least dreadful modes of conduct).

⁸ This point is altogether missed in Feinberg 1973, 86: 'For a human right to have this [absolute] character it would have to be such that no conflicts with other human rights, either of the same or another type, would be possible.'

1.1.1.6. Overtopping versus overriding

In many philosophical writings on the use of torture in extreme situations of desperation, deontological duties are said to be subject to a threshold past which they are overridden by consequentialist duties with which they conflict.⁹ The notion of overriding, invoked by such a contention, is in need of disambiguation—or, perhaps more accurately, it is in need of differentiation from the notion of overtopping.

When philosophers maintain that some duty D_1 has been overridden by a conflicting duty D_2 , they usually mean that D_1 has ceased to be operative and that the fulfilment of D_2 is therefore strongly permissible. In the eyes of most of those philosophers, overriding involves cancellation or supersession or suspension. Now, as should be evident from what has been said heretofore in this chapter, the notion of overriding—understood in the way just indicated—is unsuitable for any satisfactory account of the nature of moral conflicts. This chapter has already used ‘overtop’ and its cognates many times, and the notion of overtopping is indeed what is pertinent here in contrast with the notion of overriding. An overtopped moral obligation is not cancelled or superseded or suspended by the overtoppingly stringent moral obligation(s) with which it conflicts; rather, it retains its full force as such. It is exceeded in importance or stringency by the obligation(s) with which it clashes, but it continues to impose its requirements. Compliance with the overtoppingly stringent obligation(s) in such circumstances is weakly justified but not strongly justified. (As will be seen presently, the point calls for a reconception of the threshold envisaged by quite a few deontological philosophers.)

Just as the language of ‘overriding’ is generally best eschewed in discussions of moral conflicts, so too is the language of ‘defeasibility.’ When philosophers characterize a moral duty as defeasible, they usually mean that it is susceptible to being overridden in the sense recounted above. In the eyes of most of those philosophers, a moral duty D_1 that has been defeated by another moral duty D_2 is not merely surpassed in stringency by D_2 but is also extinguished or suspended outright by it. Given that such an understanding of clashes between moral duties does not accurately capture the persistence of each duty within any clash, we are best advised to put aside the notion of defeasibility when seeking to come to grips with these matters. That notion, like the notion of overriding, tends to foster the impression that a moral duty is deprived of its requirement-imposing force when it conflicts with a more stringent moral obligation. Because any such impression is misguided, characterizations of duties as defeated or overridden will generally be absent from this book when I am presenting my own positions rather than recounting someone else’s views.

⁹ For some examples of this threshold-deontological approach, see Fried 1978, 10; Gross 2004a, 1511–19; Gross 2004b, 230–1, 250 n4; Kadish 1989, 346; Levinson 2003, 2031–4; May 2007, 232; Moore 1997, chap. 17; Moore 2007; Nagel 1979, 56; Shue 1978, 141–3. The approach is outlined, but neither endorsed nor rejected, in Kamm 2011, 49–50; Nozick 1974, 30 n*. For some wary discussions of threshold deontology, see Coady 2011, § 6; Davis 2005, 170–3; Ginbar 2008, 24–9; Haque 2007, 629, 639–43; Harel and Sharon 2011, 848–51; Kutz 2007a, 254–6; Meisels 2008b, 174–82; Posner and Vermeule 2006, 677–80; Posner and Vermeule 2007, 40–1, 187–90; Waldron 2010a, 216–17; Ward, Johnstone, and Clucas 2009, 4; Wonnell 2011. For an often powerful critique of threshold deontology, see Alexander 2000.

Of course, the preceding paragraph is hardly suggesting that moral duties are never defeasible. Frequently, the existence of a concrete moral duty in any given set of circumstances is permanently or temporarily negated by the presence of some defeating factors. When a duty is so negated, it has come to be only *prima facie* in the epistemic sense of that phrase. Because of the presence of the defeating considerations, the putative duty does not exist as a genuine moral obligation at all. Recall, for instance, my earlier scenario of the self-defensive thrashing administered by Frank to Edward; in that scenario, the occurrence of Edward's aggressive attack has suspended Frank's normal duty to refrain from the use of violence (though Frank of course will be morally obligated not to use more force than is reasonably necessary to repel Edward's onslaught).

Nevertheless, as I have already emphasized, the sheer fact of being locked in a conflict with a more stringent moral duty never deprives any moral obligation of its status as such. Hence, when this book trains its attention on moral conflicts involving the moral prohibition on torture, it will not be broaching any ostensible moral duties that turn out to be mere appearances in particular circumstances. Because a number of philosophers and legal theorists have presumed otherwise, this book will invoke the notions of overriding and defeasibility when it summarizes their stances. However, when I elaborate my own contrasting approach to the problem of torture, those notions will disappear—for they generally obfuscate, rather than illuminate, that problem.

1.1.2. Chariness of moral conflicts

Some philosophers have been decidedly reluctant to acknowledge the occurrence of moral conflicts. The factors underlying their reluctance are varied and are in need of critical scrutiny here.

1.1.2.1. Consequentialist balancing

This chapter will shortly explore the division between deontological obligations and consequentialist obligations, in the course of delineating the general structure of morality. What should be noted at present is that a disinclination to recognize the occurrence of moral conflicts is prominent on each side of that division. However, whereas on the deontological side a disinclination of that type is confined chiefly to Kantians, it is much more widespread among consequentialists. Indeed, the very matter of conflicting moral duties has been a salient point of contention between consequentialists and deontologists during the past several decades, not least with regard to the problem of torture. While most consequentialists of sundry stripes have been intent on gainsaying the reality of moral conflicts, many deontologists outside the confines of the Kantian tradition have affirmed or presupposed the reality and importance of such conflicts.¹⁰ Deontologists are quite right to do so,

¹⁰ For some consequentialists' denials of the reality of moral conflicts, see Bobbitt 2008, 363–5; Curzer 2006, 45; Gross 2004a, 1498; Hare 1972; Hare 1981, chaps 2–3; Himma 2007, 240–1;

for—as we shall see—the persistence of deontological duties in the face of consequentialist pressures is a major source of moral conflicts.

Though some versions of consequentialism can take account of the occurrence of moral conflicts, the general consequentialist prioritization of the good over the right is largely antithetical to any clear recognition of such conflicts. Every thoroughly consequentialist doctrine takes some overarching desideratum or set of desiderata as an objective that is a touchstone for the rightness or wrongness of any mode of conduct (in accordance with the conduct's tendency to promote or impede the realization of the objective). Under such a doctrine, one's sole fundamental moral obligation is to contribute maximally to the realization of the commended objective. Hence, the supporters of such a doctrine are disposed by it to perceive any non-overtoppingly stringent moral duty as merely *prima facie* in the epistemic sense and thus as not genuinely a moral duty at all. That is, they typically believe that the factors which constitute a non-overtoppingly stringent moral duty are such as to *tend* toward the existence of a moral duty without actually giving rise to one. Actions in accordance with those factors, at the expense of factors that are more strongly promotive of some consequentialist objective, would produce the net effect of detracting from the realization of that objective. Accordingly, consequentialists are disposed to maintain that such actions are morally non-obligatory as well as morally impermissible. In the eyes of consequentialist theorists, the lone source of moral obligatoriness is the conduciveness of this or that mode of conduct to the maximal attainment of the desideratum or set of desiderata which the theorists favour. If the net effect of some mode of conduct MC would be to detract from the attainment of the specified desideratum or set of desiderata, then MC does not partake of the aforementioned source of moral obligatoriness. By the reckoning of consequentialists, then, MC is neither morally obligatory nor morally permissible. By their reckoning, nobody is under any moral duty to perform MC, and thus nobody faces any moral conflict that involves such a duty.

In short, consequentialism, with its emphasis on the balancing of considerations that respectively tend toward the existence of moral duties, is ill-equipped to deal adequately with situations of moral conflict. Of course, to say as much is hardly to reject consequentialism altogether. In regard to manifold sets of circumstances, a consequentialist emphasis on balancing is entirely appropriate; countless sets of circumstances do not pose any moral conflicts. Nevertheless, consequentialism generally obscures the occurrence of moral conflicts in the numerous situations where they do arise. Partly because of the unsatisfactoriness of consequentialism in that respect, this book's approach to the problem of torture is strongly deontological.

Paskins 1976, 143; Posner and Vermeule 2006, 676–7; Posner and Vermeule 2007, 187. For a discussion of utilitarians' denials of the reality of moral conflicts, see Sinnott-Armstrong 1988, 74–81. For some deontologists' affirmations of the importance of moral conflicts, see Ignatieff 2004, *passim*; Lukes 2005; Nagel 1979, 73–4; Walzer 1973; Williams 1965.

Unlike consequentialists, then, this book can properly address an extreme situation in which the use of torture as an interrogational technique by public-safety officials is the only means likely to prove effective in averting some cataclysmically dire crime. Consequentialists such as utilitarians would contend that the officials' resorting to the use of torture is morally permissible—strongly permissible—as well as morally obligatory (Pettit 1991, 234). Very different is the verdict delivered by this book's deontological ruminations. Although the public-safety officials are morally obligated to secure the lives of citizens against calamities that are preventable and reasonably foreseeable, and although the officials' obligations are extremely weighty when the dangers are immense, those obligations do not cancel or suspend any moral duties with which they conflict. In particular, they never cancel or suspend the officials' moral obligation to refrain from the use of torture. On a prodigiously rare occasion when a moral obligation of the latter type is overtopped in ethical importance by some countervailing moral obligations of officials to protect citizens against extraordinary perils, the officials' duty to refrain from the use of torture persists as such. Albeit the morally optimal course of conduct for the officials in the terrible circumstances is to act athwart that duty, they will indeed be acting athwart it rather than acting permissibly. Sometimes a morally vital course of conduct is morally impermissible. Though anybody faced with such a moral clash will have acted correctly if he fulfils the weightier of the two conflicting moral obligations, he will have breached the less important moral duty and will thus have incurred a further moral duty to remedy the situation in some way. Were he acting permissibly by fulfilling the weightier duty, he would not incur any remedial obligation; yet, precisely because he will not in fact be acting permissibly (even though he is acting optimally), such an obligation will indeed be incurred.

Admittedly, consequentialist theorists such as utilitarians can allow that some measures which *resemble* genuine remedies should be undertaken by legal-governmental officials after torture has been employed in a utility-maximizing fashion. If the use of such torture has somehow become known, and if numerous people in the officials' society do not themselves adhere to utilitarian ways of thinking, then the application of a remedy-resembling measure will most likely tend toward the maximization of utility through its alleviation of people's distress. Such a measure might be adopted with full sincerity by the officials who implement it, but utilitarian theorists have to perceive it as a simulacrum of a genuine remedy. A veritable remedy constitutes an acknowledgement that something wrong has been done. That is, it constitutes an acknowledgement that some moral duty has been breached.¹¹ Utilitarian theorists are obliged by their doctrine to deny that any wrong whatsoever has been committed when legal-governmental officials resort to torture in desperate circumstances where such a course of action is foreseeably

¹¹ Thus, for example, a remedial obligation is different from an obligation to make a payment for something that has been legitimately acquired (such as a meal consumed by a customer in a restaurant or a plot of land acquired by a government through its power of eminent domain). A payment for something legitimately acquired does not *remedy* a wrong that has occurred; rather, it *averts* the occurrence of a wrong. See Kramer 2004, 286–7.

utility-maximizing. Those theorists are therefore obliged by their doctrine to deny that the remedy-resembling step broached above is genuinely remedial. Given that there is no wrong to be rectified, the aforementioned step does not play any rectificatory role. That remedy-resembling measure will be applauded by utilitarian theorists if it tends toward the maximization of utility, but they will not view it as a way of correcting a wrong. In their eyes, its only role is consolatory. It mollifies people who erroneously think that they or others have been treated impermissibly. A utilitarian theorist will maintain that, if legal-governmental officials are morally obligated to take any remedy-resembling step in the circumstances envisaged here, their obligation is wholly derivative of their general moral duty to make people feel happy (in this case, by catering to people's misconceptions).¹² Such a theorist will contend that the obligation of the officials has nothing to do with any need for the righting of a wrong, since no wrong has occurred.

Utilitarianism's gainsaying of the reality of moral conflicts is to be rejected, for it entails the conclusion that absolutely any policies—the execution of innocent people, the most brutal varieties of torture, genocide—are morally permissible so long as the utility-promotive considerations that support the policies are strong enough. We should recognize instead that someone faced with a moral conflict will have to act impermissibly regardless of what he does. If legal-governmental officials are ever confronted with a predicament in which they elect to breach their moral duty-not-to-avail-themselves-of-torture in order to comply with some surpassingly important moral duties owed to the general public, they will be justified only in the sense that they will be pursuing the morally best course of action. They will not be justified in the sense of behaving permissibly. On the contrary, because any use of interrogational torture is abidingly impermissible, the officials will trigger stringent remedial obligations if they ever have recourse to such a tactic. Those obligations will be genuinely remedial, rather than a sop to people's delusions.

1.1.2.2. The objectivity of morality

Bernard Williams helped to bring the topic of moral conflicts to prominence in Anglo-American philosophy through his publication of a couple of incisive and influential essays in the mid-1960s (Williams 1965; 1966). Nonetheless, commendable though his contribution was, he muddied the waters in one significant respect. He suggested that the actuality of moral conflicts is somehow problematic for moral realism; in other words, it is somehow problematic for the thesis that morality is ontologically and epistemically and semantically objective.¹³ Were Williams correct

¹² If the remedy-resembling step is a simulacrum of compensation, then the utilitarians can favour making it mandatory on the ground that requiring such an expenditure will internalize the costs of governmental operations and will thereby help to ensure that those operations in the future do not go ahead unless their expected drawbacks are significantly exceeded by their expected benefits. Many theorists in the law-and-economics movement ascribe a non-remedial role of this type to awards of damages in tort law.

¹³ I have elsewhere defended at length a minimalist non-naturalistic version of moral realism. See Kramer 2009a.

on that point, there would obviously be a strong reason for proponents of moral realism (including me) to deny that moral conflicts do in fact occur.

There is no need here for a lengthy engagement with Williams's arguments about moral realism, since—unlike his far more perceptive exploration of the general structure of moral conflicts—those arguments have been discredited for quite some time.¹⁴ Moreover, as will be seen, his main way of drawing an anti-realist inference from his reflections on moral conflicts is a version of an unsound line of reasoning that will be probed in § 1.1.2.3. Thus, a terse rebuttal of his anti-realist pronouncements will be sufficient here.

Williams's central point in his challenge to moral realism is articulated in the following passage from his 1965 essay:

It seems to me a fundamental criticism of many ethical theories that their accounts of moral conflict and its resolution do not do justice to [the persistence of conflicting moral duties]: basically because they eliminate from the scene the 'ought' that is not acted upon. A structure appropriate to conflicts of belief is projected on to the moral case; one by which the conflict is basically adventitious, and a resolution of it disembarrasses one of a mistaken view which for a while confused the situation. Such an approach must be inherent in purely cognitive accounts of the matter; since it is just a question of which of two conflicting 'ought' statements is true, and they cannot both be true, to decide correctly for one of them must be to rid of error with respect to the other. (Williams's 1965, 113)

On the one hand, Williams was undoubtedly correct when he wrote that a number of ethical theories have failed to do justice to the persistence of moral duties in moral conflicts. On the other hand, he went badly astray in submitting that moral realists are at any disadvantage in coming to grips with the nature of moral conflicts. His error becomes apparent in the final sentence of this quotation, which asserts that two conflicting 'ought' statements cannot both be true. In fact, both the statement 'Some person F at some time t is under a moral duty-to- ϕ ' and the statement 'Some person P at some time t is under a moral duty-not-to- ϕ ' can be true. In the final sentence of this quoted passage, Williams was presupposing that 'P at time t is under a moral duty-not-to- ϕ ' is equivalent to 'It is not the case that P at time t is under a moral duty-to- ϕ .' His presupposition is untenable, however, as it conflates internal and external negation. Although the proposition 'It is not the case that P at time t is under a moral duty-to- ϕ ' and the proposition 'P at time t is under a moral duty-to- ϕ ' cannot both be true (and cannot both be false), the proposition 'P at time t is under a moral duty-to- ϕ ' and the proposition 'P at time t is under a moral duty-not-to- ϕ ' can perfectly well both be true (and can both be false).

What is so peculiar about the misstep by Williams is that elsewhere in his 1965 essay he himself highlighted the distinction which he elided in the passage above. As Philippa Foot aptly commented: 'The strange thing about what Williams

¹⁴ Of key importance in discrediting those anti-realist arguments was Foot 1983. For another excellent rebuttal of those arguments, see Brink 1994, 242–6.

wrote in “Ethical Consistency” is that a great deal of it seems designed to show exactly... that moral conflict does not imply “contradiction.” It is as if he himself showed the cognitivist how to avoid the very error he thinks the cognitivist must make’ (Foot 1983, 391).

1.1.2.3. *Logical incoherence*

Some of the philosophers who deny the possibility of moral conflicts are worried by the thought that such conflicts would partake of logical incoherence.¹⁵ We have just seen that a thought along those lines is central to Williams’s indictment of moral realism. Williams usually knew better, but some other philosophers have been much more firmly resistant to the notion that moral conflicts are logically possible. For example, Richard Hare, who was staunchly resistant to that notion throughout his career, sometimes conflated internal and external negation in much the same manner as Williams. In an essay published in 1989, he wrote as follows: ‘If I say “I ought, but there is someone else in exactly the same circumstances, doing it to someone who is just like the person I should be doing it to, but he ought not to do it,” then logical eyebrows will be raised; it is *logically inconsistent* to say, of two exactly similar people in exactly similar situations, that the first ought to do something and the second ought not’ (Hare 1989, 179, emphasis in original). Seeking to expose a logical inconsistency, Hare placed the negation in the wrong position. Whereas ‘I ought to ϕ ’ logically contradicts ‘It is not the case that I ought to ϕ ,’ it is logically consistent with ‘I ought not to ϕ .’ In other words, ‘not ought’ rather than ‘ought not’ is the contradictory of ‘ought’. Given that the structure of a moral conflict comprises ‘I ought to ϕ ’ and ‘I ought not to ϕ ,’ and given that no role is played in any moral conflict by ‘It is not the case that I ought to ϕ ,’ there are no logical inconsistencies in such a conflict. As a logical matter, the occurrence of moral conflicts is entirely unproblematic.

Quite a few other philosophers have committed mistakes similar to that of Hare (Kramer 2009a, 335 n17). One of the central principles of standard deontic logic, the so-called Permissibility Theorem, tends to cloud reflection on these matters.¹⁶ The formal rendering of that theorem is as follows:

$$(\forall x)(Ox \rightarrow Px)$$

¹⁵ This worry surfaces fleetingly in Coady 2011, § 3–4; Nagel 1979, 59–60, 74 n12; Wisniewski 2010, 62–3. It is much more prominent in Curzer 2006. It is discussed and rejected in Meisels 2008a; Meisels 2008b, chap. 7.

¹⁶ Some sophisticated philosophers have endorsed the Permissibility Theorem. See, for example, Conce 1982; Feinberg 1980, 235, 237; Hill 1996, 177; Hughes and Cresswell 1968, 43; McConnell 2010, § 4; Simister 2008, 292 n10; Steiner 1998, 233, 268 n55; Vallentyne 1987, 119–20; Vallentyne 1989. For a powerful critique, see Sinnott-Armstrong 1988, 156–61. See also Sinnott-Armstrong 1996, 52. Though David Brink (1994, 235–6) appears to endorse the contrapositive of the Permissibility Theorem, he in fact ultimately endorses only the version of that theorem which I myself uphold.

What this formulation states in effect is that the obligatoriness of any mode of conduct entails the permissibility thereof. By contraposition, then, the impermissibility of any mode of conduct entails the non-obligatoriness thereof. Now, although there is a way of construing the Permissibility Theorem which renders it necessarily true, that one truth-conferring interpretation of the theorem is fully consistent with the possibility and actuality of moral conflicts. Each of the three other eligible interpretations of the theorem falsifies it. On the whole, the Permissibility Theorem has impeded rigorous thinking about the structure of deontic relations.

We are told by the Permissibility Theorem that everything obligatory is permissible. If we construe 'obligatory' as 'overtoppingly obligatory', and if we construe 'permissible' as 'at least weakly permissible', then the Permissibility Theorem is true. It is indeed the case that every overtoppingly obligatory course of conduct is at least weakly permissible; if some such course of conduct were not at least weakly permissible, then a person would be under an overtoppingly stringent moral duty-to-do- q and an overtoppingly stringent moral duty-not-to-do- q . Given that an overtoppingly stringent moral duty exceeds in importance any moral duties that conflict with it, each of the two duties just mentioned would be morally more important than the other. Such a state of affairs is impossible, since any coherent relation of superiority (such as 'more important than') is strictly asymmetrical. Ergo, every overtoppingly obligatory course of conduct is indeed at least weakly permissible.

As has been remarked, this one truth-conferring interpretation of the Permissibility Theorem renders it straightforwardly compatible with the existence of moral conflicts. In any moral conflict, where someone is under a moral duty-to- ϕ and simultaneously under a moral duty-not-to- ϕ , no more than one of those obligations is overtoppingly stringent. Accordingly, every such conflict is consistent with the Permissibility Theorem as it has just been glossed. Everyone who recognizes the coherent possibility of moral conflicts should also recognize that the Permissibility Theorem as it has just been glossed is true.

Construed in any other way, however, the Permissibility Theorem is false. It is not the case, for example, that every overtoppingly obligatory course of action is strongly permissible. One's moral duty to aid a badly injured pedestrian might well be overtoppingly stringent in a given context, but the fulfilment of that overtoppingly stringent duty might preclude the fulfilment of one's less important moral duty to keep one's promise to meet somebody else for lunch. In that event, one's going to the aid of the pedestrian is only weakly permissible rather than strongly permissible. Innumerable similar examples could be adduced.

Even more plainly, it is not the case that every non-overtoppingly obligatory course of action is strongly permissible. Indeed, no such course of action is strongly permissible. Furthermore, given that some non-overtoppingly stringent moral duties conflict with overtoppingly stringent moral duties, it is not even the case that every non-overtoppingly obligatory course of action is weakly permissible. Many such courses of action are impermissible in every sense. Hence, only under one interpretation is the Permissibility Theorem correct; under the three other available interpretations, it is false.

Of central importance here, again, is that the lone truth-conferring interpretation of the Permissibility Theorem is consistent in every respect with my emphasis on the possibility and actuality of moral conflicts. That rendering of the theorem is a wholly unexceptionable encapsulation of the logic of the 'greater than' relationship. Since every coherent moral conflict naturally conforms to that logic, the Permissibility Theorem poses no obstacles whatsoever to the occurrence of moral conflicts. However, because that theorem is so susceptible to being construed in any of the three ways that render it false, it tends to be highly misleading. It lends itself too readily to egregious misapprehensions. We are therefore best advised to leave it aside, even while recognizing that it is completely unproblematic when it is understood correctly.

1.1.2.4. Moral conflicts and action-guidance

Moral theories and principles are supposed to provide answers to questions about appropriate courses of conduct in multitudinous sets of circumstances. Thus, one complaint sometimes voiced about the spectre of moral conflicts is that we would have no grounds for deciding what to do when confronted with clashing obligations (Dworkin 2011, 90; McConnell 2010, §§ 4 and 7). In other words, moral conflicts are thought to leave us bereft of action-guidance. Any anxiety along those lines is largely baseless. In most moral conflicts, the clashing duties are of unequal stringency; hence, the uniquely correct response to such a conflict is to fulfil the more stringent duty. Of course, as should be evident from my earlier discussion of the adjective 'right', the term 'correct' here does not denote permissibility. Within a moral conflict, no morally permissible course of conduct is available. Nonetheless, although the uniquely correct mode of conduct in such circumstances is itself morally wrong—and although the adoption of that course of conduct will therefore give rise to remedial duties—it is indeed uniquely correct in that the non-adoption of it would be an even more serious wrong. Accordingly, there is no lack of action-guidance in any such moral conflict. Morality determinately prescribes the mode of conduct that is to be undertaken as the best way of dealing with the quandary which such a moral conflict presents.

Furthermore, even in a rare moral conflict where the clashing duties are evenly balanced in their stringency or are incommensurably counterpoised, the conclusion follows that the fulfilment of either duty will amount to a breach of the other and will thus impose remedial obligations on anyone who has committed the breach.¹⁷ Hence, although neither the fulfilment of the duty-to- ϕ nor the fulfilment of the duty-not-to- ϕ in such a conflict is a morally better course of conduct than the other, we are not wholly devoid of guidance in a situation of this sort. Anyone aware of the nature of the situation as a moral conflict has grounds for knowing that, regardless of whether he complies with the former duty or with the latter, he will have incurred a further moral obligation to remedy the wrong that he has thereby done.

¹⁷ This point is rightly emphasized in Nussbaum 2000, 1009.

1.1.2.5. *Concerns about fairness or excessive onerousness*

In short, the problem of action-guidance is no more powerful as a basis for some philosophers' chariness of moral conflicts than are any of the other considerations that we have pondered so far. A rather different concern appears to underlie a remark by Michael Moore in which he explains why he hopes to show that moral conflicts are very uncommon. On the one hand, Moore does briefly acknowledge the potential for conflicts between deontological duties, and he allows that '[i]t may not be, as Kant famously proclaimed, that a conflict of such obligations is literally "inconceivable"' (2007, 37). On the other hand, he declares that 'it would be unfortunate for us in the extreme if morality often confronted us with choices where we will be "damned if we do and damned if we don't." The distinctions we shall examine hold out the possibility of so limiting our stringent obligations as to minimize or even eliminate such situations of moral conflict' (2007, 37–8). Indeed, the very scenarios with which he grants the possibility of conflicts between deontological obligations—scenarios similar to some of those in Kamm 2007, 27–9, 252, 268–9—are so far-fetched as to make manifest his view that any genuine conflicts in morality are extremely rare.¹⁸

Moore's ambition to establish the infrequency of moral conflicts is focused not on the problem of action-guidance but instead on the onerousness or unfairness of a world in which someone often finds that every mode of conduct open to her is morally wrong. Such a worry is pertinent, of course, but Moore draws an inapt conclusion from it. Instead of trying to expound the general structure of morality in a manner that whisks most moral conflicts out of sight, we should quite frequently seek to act in ways—and to arrange our institutions in ways—that will reduce the incidence of such conflicts. Acting in conformity to a practically oriented conclusion of that sort is the best means of allaying the anxiety felt by Moore and others about the prospect of unavoidable wrongness. To act in such a fashion, one needs to be alert to the possibility and actuality of moral conflicts in a diversity of settings.

In other words, the consternation engendered by the spectre of unavoidable wrongness should prompt us to do the opposite of what Moore recommends. Far from trying to delineate the contours of morality in a manner that will obscure the emergence of moral conflicts, we should be seeking to grasp those contours with keen sensitivity to the likelihood of such conflicts. Only thus can we informedly fix upon the practical steps that are best suited to avert predicaments of unavoidable wrongness (insofar as they can and should be averted).

¹⁸ For instance, Moore propounds the following scenario (2007, 96): 'I have begun a boulder rolling down a hill to kill five old enemies of mine, but... I repent. Now, however, the only way to stop the boulder is with the corpulent body of one bystander: may I throw him before the boulder, killing him but preventing the deaths of the five? Doing so makes me a killer, but it prevents me from being a killer five times over.'

1.2. The Deontology/Consequentialism Distinction and the General Structure of Morality

As has been remarked, this book's approach to torture is embedded in a general account of the nature of morality. Of key importance to that account is my emphasis on moral conflicts, which we have just been considering. Likewise of key importance is the strongly deontological tenor of this book's outlook. Having already glanced at the distinction between deontology and consequentialism in my ruminations on moral conflicts, we should now examine it more directly and sustainedly. Thereafter, this chapter will conclude by adumbrating the structure of morality. Naturally, both the matter of moral conflicts and the deontology/consequentialism dichotomy will figure saliently in the adumbration.

1.2.1. The division between deontology and consequentialism

My outline of the general structure of morality, along with the rest of this book, will refer prominently to deontological prohibitions and deontological permissions (or prerogatives). Some careful attention here to the deontology/consequentialism distinction will help to elucidate those references. We should mull over two main ways in which that distinction can be drawn.

1.2.1.1. *Intrinsic moral statuses*

One tack for differentiating between deontology and consequentialism is to concentrate on the moral statuses of modes of conduct. If a course of conduct on the part of any person *P* is covered by a deontological permission, then *P*'s engaging in that conduct is not wrong in any respect—regardless of the consequences that it causes or is likely to cause. Within the ambit of the deontological permission, *P*'s undertaking of the specified conduct does not breach any moral duties. Hence, within that ambit, his behaviour is strongly permissible irrespective of the probable consequences of its occurrence. Suppose, for example, that *P* is deontologically at liberty to expend a certain portion of his income on his hobby of collecting stamps. Suppose further that any number of consequentialist objectives (such as the maximization of human happiness or the promotion of equal economic opportunities) would be advanced if *P* were instead to expend that portion of his income on donations to charities. Notwithstanding that his pursuit of his philatelic hobby will bring about a worse state of affairs than he is capable of bringing about through alternative uses of his resources, *P* is not committing any wrongs so long as he is acting within the scope of the deontological prerogative which encompasses that pursuit.

Just as the permissibility of *P*'s expenditure of funds on his hobby is unaffected (within broad limits) by the consequences of his using his income in that fashion, so too the wrongness of certain modes of conduct is not removed by the benignity of the consequences that are likely to flow from the occurrence of any such modes

of conduct. If P is under a deontological duty to ϕ , then his not ϕ -ing is wrong regardless of how valuable the resultant state of affairs might be. Every deontological duty is at least weakly absolute (in the sense introduced earlier). That is, every deontological duty forbids some course of conduct even if one's undertaking of that course of conduct in particular circumstances would yield much better consequences than one's refraining therefrom. A deontologically prohibited type of conduct is wrong always and everywhere.

Whereas deontologists ascribe consequence-independent moral bearings to various actions and omissions, consequentialists of course maintain that the moral character of any type or instance of conduct is fully determined by the probable consequences thereof. In the eyes of the latter theorists, no type or instance of conduct is ever endowed with any inherent moral status. Instead, every action or omission derives its moral status from the effects with which it is associated. Accordingly, as John Finnis has aptly remarked, a proponent of any consequentialist doctrine 'holds himself ready to do *anything*' (Finnis 1980, 121, emphasis in original). For the consequentialist, there are no principled limits on the range of actions that can legitimately be undertaken in sundry circumstances. The breadth or narrowness of that range will depend entirely on the results that are likely to follow from each of the multitudinous modes of conduct that might be adopted by people in any number of contexts. No mode of conduct, however abhorrent it may be, is unconditionally disallowed by a consequentialist theorist.

1.2.1.2. *Agent-neutrality versus agent-centredness*

One way of differentiating between deontology and consequentialism, then, addresses the question whether the moral statuses of some patterns of behaviour are intrinsic or extrinsic. Another way—which has become especially prominent during the past few decades and which figures conspicuously in some of the philosophical literature on torture—pertains to the manner in which our fundamental moral obligations present their demands. Whereas deontologists contend that those demands are agent-centred,¹⁹ consequentialists maintain that they are agent-neutral. In other words, deontological duties are such that they present each agent with reasons specifically for that agent to conduct himself in certain ways. The reasons do not pertain to goals by reference to which every agent is enjoined to produce maximally valuable or minimally disvaluable states of affairs. For example, a deontological duty to refrain from the crime of murder is not a duty to pursue an overarching goal such as the minimization of the number of murders in a society. Rather, it is a duty that requires each person to refrain from committing any murders even when his perpetration of such crimes would reduce their overall incidence or would otherwise be highly valuable. One's fulfilment of a deontological obligation

¹⁹ See, for example, Nagel 1979, 132–4; Scheffler 1985; Darwall 1986. The epithet 'agent-relative' is also common in the relevant literature—owing largely to Parfit 1984, 143 and Nagel 1986, 152—but I eschew it in order to make fully clear that a deontological conception of moral duties has nothing to do with moral relativism.

consists not in one's venturing alongside other people to maximize the realization of some objective, but instead in one's compliance with a strict prohibition regardless of how other people might respond.

By contrast, consequentialist obligations require each person to strive alongside other people to maximize the realization of some goal or array of goals (which might be highly pluralistic). A consequentialist obligation is agent-neutral in that it presents everyone with reasons-for-action that do not single him or her out as a particular person distinct from everybody else. In that respect, the reasons flowing from any consequentialist duty are the same for everyone—though of course the detailed specificities of what is required of each person will depend on his or her aptitudes and circumstances. In regard to any desideratum or set of desiderata upheld by such a duty, a maximizing quest is prescribed. From a consequentialist perspective, the attainment of the most valuable state of affairs possible is always better than the attainment of any less valuable state of affairs and is therefore both obligatory and strongly permissible. Everyone alike is morally obligated and morally at liberty to contribute to the achievement of that maximally valuable state of affairs, even if someone's contribution will involve his acting at odds with the content of a deontological principle. (Of course, although any consequentialist theory maintains that the maximization of some desideratum or set of desiderata is morally obligatory and permissible, not every such theory recommends that individuals adopt maximizing outlooks when arriving at decisions. Many consequentialists instead hold that the maximization of a favoured desideratum or set of desiderata is best pursued indirectly through the focusing of individuals' choices and deliberations on other concerns. For example, the adoption of a satisficing mentality by individuals in relation to some desideratum might be the best way for them to maximize their attainment of a more complex balance of desiderata.)

1.2.1.3. Absolute prohibitions

Such, then, are the two ways of distinguishing between deontology and consequentialism. Now, as is plain, those two ways are not identical. One of them addresses the question whether the basic moral statuses of certain modes of conduct are intrinsic or extrinsic,²⁰ whereas the other addresses the question whether the reasons presented by fundamental moral obligations are agent-centred or agent-neutral. Nonetheless, although the two contrasts are not exactly the same, they are not really separate; the latter is a corollary of the former. If the basic moral status of a mode of conduct is consequence-independent, then—in the manner captured by the notion of agent-centredness—that basic moral status is impervious to any consequentialist balancing that would deem the tokens of an abhorrent act-type AT to be permissible whenever those tokens can lower the overall incidence of AT. Imperviousness to such balancing is a key aspect of the general consequence-independence of an act-type's wrongness.

²⁰ The phrase 'basic moral status' refers to the permissibility or impermissibility of each of the modes of conduct.

The distinction between deontology and consequentialism has a crucial bearing on the matter of the absoluteness (the weak absoluteness) of some major prohibitions. On the one hand, as this chapter has already observed, not all deontologists are absolutists. Some threshold deontologists believe that the consequence-independent impermissibility of a mode of conduct such as torture is operative only up to a certain point—namely, a point where the consequences of forgoing that mode of conduct would be truly appalling. On the other hand, many deontologists (including me) are indeed absolutists. Such deontologists affirm that certain major moral prohibitions are binding always and everywhere irrespective of the consequences of compliance therewith, and specifically that the prohibitions are binding even when departures from them will serve to lower greatly the overall incidence of such departures. By contrast, consequentialists deny that any prohibitions are absolute (apart from a general prohibition on declining to contribute to the realization of some overriding desideratum or set of desiderata in a maximally promotive manner). In particular, consequentialists deny that any moral prohibition remains binding when a deviation from it will serve to heighten the overall level of conformity to that prohibition.

While the distinction between deontology and consequentialism will loom large throughout this book, we shall return to it most sustainedly in Chapter 4—a chapter that seeks to confirm the absoluteness of the moral prohibition on torture. As will become apparent there, some philosophers have mounted powerful challenges to the agent-centredness of deontological prohibitions. Hence, one part of my task in defending the absoluteness of the moral prohibition on torture will be to reaffirm the rationality of deontological agent-centredness in the face of those challenges. A vindication of deontology is essential for a vindication of the thesis that torture is always and everywhere wrong.

These present remarks on the deontology/consequentialism division should close with a terse caveat that can help to forestall confusion hereafter. Chapter 3's discussion of the wrongness of torture will lay emphasis on the distinction between a victim-focused perspective and a perpetrator-focused perspective for the justification of a moral prohibition. That is, it distinguishes between (1) justifications that concentrate on the interests or inviolability of the potential victims of torture and (2) justifications that concentrate on the moral integrity of the potential perpetrators of torture.²¹ Although a rationale of the first type is perfectly consistent with a rationale of the second type, neither entails the other; either can be invoked to the exclusion of the other, even though they can also be combined. Chapter 3 will argue that a victim-focused rationale is insufficient on its own to support the thesis that torture is always and everywhere wrong. Notwithstanding that the interests and inviolability of potential victims of torture are decisive justificatory factors in most contexts, a victim-focused approach uncombined with a perpetrator-focused approach will fall short of establishing that the moral bar to the permissibility of

²¹ This distinction between a victim-focused justification and a perpetrator-focused justification is similar to the distinction which I have drawn elsewhere between right-focused justifications and duty-focused justifications; see Kramer 1998, 35–41.

torture is absolute. A satisfactory engagement with the problem of torture has to concentrate on the moral integrity of people who might resort to the deliberate infliction of excruciating pain.

The main reasons for the inadequacy of a purely victim-focused perspective will come to the fore in Chapter 3. For the moment, we should simply note that the victim-focused/perpetrator-focused dichotomy is quite different from the agent-neutrality/agent-centredness contrast which we have briefly explored above. Whereas the distinction between a victim-focused justification and a perpetrator-focused justification pertains to the substantive basis for a moral prohibition on the use of torture (or the substantive basis for some other moral constraint), the agent-neutral/agent-centred dichotomy bears on the question whether any such prohibition is susceptible or insusceptible to consequentialist trade-offs. To be sure, the moral considerations that account for the forbiddenness of torture might well be the same as the moral considerations that account for the fact that the forbiddenness of torture is not subject to consequentialist exceptions. However, any such convergence of considerations is something that has to be shown through moral argumentation, rather than something that can safely be taken for granted through the conflation of the two issues or sets of issues that have just been disentangled here. Even if the factors that underlie the wrongness of torture are also the factors that make the moral prohibition on torture impervious to any consequentialist balancing, the wrongness and the imperviousness are not exactly the same thing. Questions about those matters have to be addressed separately, even if the answers to the questions exhibit close affinities with each other.²²

1.2.2. The general structure of morality

Keeping in view the foregoing exposition of the deontology/consequentialism distinction, we can now turn to the general structure of morality—in which that distinction will figure centrally. To highlight the distinctiveness of this book's conception of that general structure, the present section of the chapter will juxtapose Moore's account of morality with my own account. Moore's presentations of his understanding of morality have been helpfully expansive and sophisticated, and his understanding is in accordance with the views of many other philosophers who embrace a threshold-deontological position. Because his perspective (a quite widely shared perspective) is at variance with this book's version of threshold deontology in several major respects, his delineation of his moral framework is an excellent foil for the elaboration of my competing map of the moral realm.

²² The differences between the agent-centred/agent-neutral dichotomy and the perpetrator-focused/victim-focused dichotomy are implicitly recognized in Gewirth 1981, 14; and Hill 1991, 81. They are overtly recognized in Kamm 2007, chap. 8. They seem to be overlooked in Gross 2004a, 1494 n48; and Moore 1997, 705 n89. They are definitely overlooked in Waldron 1989, 505 n5. Some related confusion arises when the agent-centred/agent-neutral distinction is perceived as a variant of the act/omission distinction. For an example of such confusion, see Seidman 2005, 889–91.

Moreover, as will become evident later in the book, these divergent outlooks on very abstract issues of morality ramify into divergent assessments of the concrete moral matter of interrogational torture. Dissimilarities between my conclusions about interrogational torture and Moore's conclusions relate not just to the moral bearings of such torture but also to the legal consequences that should ensue if torturous methods of interrogation are ever employed in extreme situations. Underlying the contrarities between our proposals concerning those legal consequences is the matter of moral conflicts that has loomed so large already in this chapter.

To be sure, the connections between the abstract issues about moral conflicts and the concrete issues about the proper legal treatment of torture are not logically necessary. Those connections are substantively moral,²³ and the substantive moral values at stake can be explicated plausibly in a number of directions. Nonetheless, as the trajectory of this book will suggest, the paths between the abstract heights and the concrete thickets of moral controversy are sometimes quite straightforward. Theses propounded on the abstract heights can be enrichingly fleshed out and supported by the concrete moral reasoning for which these theses furnish indispensable guidance.

1.2.2.1. Moore's tripartite account

Moore maintains that the general structure of morality is tripartite. At the first level is a background of consequentialist reasons for each person to seek to bring about valuable states of affairs. Those reasons are always present in varying degrees of weightiness, though they do not always impose obligations. Whereas some consequentialist theories posit a single fundamental desideratum (such as maximal human happiness) that is the ground of the value of everything else, Moore's conception of the good is pluralistic. Everyone has reasons for promoting sundry types of valuable states of affairs.

As has been remarked in the preceding paragraph, consequentialist reasons-for-action do not always impose obligations. The absence of obligatoriness is quite often due to the normative effects of the second of the three levels in the structure of morality that Moore envisages. At that second tier are deontological permissions (or prerogatives) and deontological duties. Deontological permissions strongly entitle people to conduct themselves in ways that are suboptimal from any consequentialist viewpoint, while deontological prohibitions forbid people to act in certain ways even when their so acting would be optimal from any consequentialist viewpoint. Thus, the second tier in Moore's triadic account of morality is an expression of his resistance to the notion that consequentialist trade-offs exhaust the moral domain.

Nevertheless, the third tier in his account makes clear that his moral outlook is not steadfastly deontological.²⁴ As a threshold deontologist, Moore maintains that

²³ Of course, as I have argued elsewhere, the correct basic principles of morality and many correct derivative principles of morality are themselves necessary. See Kramer 2009a, 157–61. My point here is simply that the necessity is moral rather than logical.

²⁴ Chapter 4 will address another respect in which Moore's adherence to deontology is somewhat tenuous: namely, his uneasiness about the rationality of the ways in which deontological duties check

the persistence of deontological duties in response to consequentialist buffeting is not unlimitedly tenacious. Past some unspecifiable threshold as the consequences of adhering to a deontological duty *D* become more and more dire, *D* ceases to be binding. In a context where the threshold of catastrophic moral horror has been passed, calamity-averting actions at odds with *D*—which, for the duration of the extreme emergency, is only an erstwhile obligation—are morally permissible. For Moore, then, the third layer in the realm of morality is made up of situations of grave peril where the urgency and immense weightiness of consequentialist requirements override (and not merely overtop) the sway of deontological duties.

1.2.2.2. An alternative account of morality

My alternative theory of the general structure of morality overlaps with Moore's, but is much more robustly deontological. My theory, like his, recognizes that each person is confronted with an array of consequentialist reasons-for-action. Those reasons-for-action, which are of many different degrees of strength, can constitute moral obligations of varying degrees of stringency. Insofar as endeavours to further some genuine desiderata do not involve any contraventions of deontological constraints, people beyond the bounds of their deontological prerogatives will generally be morally obligated to participate in those endeavours (in specific ways that will depend on their aptitudes and circumstances).

In addition to being confronted with an array of consequentialist reasons-for-action, every person is faced with an array of deontological permissions and prohibitions. Of most importance here are conflictual relationships between those reasons-for-action on the one hand and those permissions and prohibitions on the other. When consequentialist considerations clash with the demands of a deontological prohibition, and when the specified considerations are not very weighty, they do not constitute any moral obligations. Instead, the matter to which they pertain is covered by a deontological prerogative, which is thus coupled with the deontological prohibition. In other words, conformity to the deontological prohibition is strongly permissible even though such conformity is inconsistent with what a consequentialist calculation prescribes.

When the consequentialist factors that clash with a deontological prohibition are very weighty, those factors are constitutive of a countervailing moral obligation. In such a situation, accordingly, the prohibition is no longer coupled with a deontological permission. Instead, it is locked in a conflict with a diametrically opposed moral requirement; the deontological duty-not-to- ϕ is countervailed by a consequentialist duty-to- ϕ . Still, that latter duty is exceeded by the former in stringency. If any

the maximization of desiderata (Moore 1997, 705; Alexander and Moore 2007, § 2.2.2). Moore is worried about the powerful challenge by Samuel Scheffler (1985) to the rationality of agent-centred restrictions—a challenge to which I have referred briefly in § 1.2.1.3. Chapter 4 will seek to rebut Scheffler's arguments, and will present two complementary rationales for the agent-centredness of the aforementioned restrictions.

deontological duty D is strongly absolute, then in every possible world D is always and everywhere more stringent than any competing obligations. If a deontological duty is only weakly absolute, its stringency can be exceeded by that of a counter-vailing moral obligation in a situation of extreme urgency and desperation. Exactly how dire the situation would have to be for the overtopping of the deontological duty's normative importance is something that depends on the specifics of the duty and on the extent of any requisite contravention.

What should be emphasized here, in line with my earlier remarks on the persistence of any moral duties that are locked in a moral conflict, is that an overtopped deontological obligation retains all its normative force despite being surpassed in stringency by a formidably weighty consequentialist requirement that conflicts with it. Its demands are uncancelled and indeed unimpaired by the situation of moral conflict in which it obtains. Admittedly, the remedial duties arising from a breach of the obligation will be considerably less onerous than the remedial duties that would have arisen if the breach had occurred in the absence of any conflict. The gravity of a contravention is substantially mitigated when the contravention is undertaken for the sake of fulfilling an overtoppingly stringent moral requirement. All the same, the very fact that remedial duties will indeed accrue in the wake of a morally optimal course of action is attributable to the undiminished force of the moral obligation that has been transgressed. The status of that deontological obligation as such is unaltered by the presence of an even more important moral demand that calls—in an extreme context—for the obligation to go unfulfilled.

1.2.2.3. A pithy conclusion: some differences between the accounts of morality

Moore's account and my own account of the general structure of morality are both threshold-deontological in character, but our conceptions of the threshold differ markedly. Whereas Moore takes it to be a threshold of moral permissibility, I take it to be only a threshold of moral optimality. This key difference is directly connected to his discounting of the significance of moral conflicts and to my highlighting of that significance. To be sure, this major divergence between our approaches is accompanied by an array of similarities between them. For example, although we disagree about the nature of the deontological threshold, we concur that the location of that threshold is indeterminate.²⁵ Nevertheless, the affinities between our understandings of morality are overshadowed by the dissimilarities.

As has been indicated, those dissimilarities carry over from the abstract reaches of moral philosophy into Moore's and my divergent assessments of the moral bearings of torture—and into our even more concrete recommendations concerning the legal consequences that should ensue if torture is ever employed. Precisely

²⁵ Thus, I believe that Moore is not vulnerable to several of the criticisms of him in Alexander 2000. Though the remaining criticisms in Alexander's trenchant article are telling against Moore, they are readily accommodated by my own deontological conception of morality.

because Moore supposes that the threshold in his deontological account of morality is a threshold of permissibility as well as of optimality, he believes that no wrong whatsoever is committed when interrogational torture is employed to thwart a cataclysmically destructive act of terrorism or criminality. No moral conflict is involved, or so Moore contends. Because no moral wrong has been done, legal sanctions levied for the use of interrogational torture in the specified circumstances would be inappropriate.²⁶ Hence, in combination with his legal-moralist theory of the role of criminal law, the disinclination of Moore to acknowledge the frequency of moral conflicts has impelled him quite smoothly toward his position on the proper legal status of torture.

Likewise, of course, my insistent acknowledgement of the frequency of moral conflicts is a key to my position on the legal consequences that should follow when interrogational torture has been employed. Even in an imaginably dreadful situation of urgency where the use by officials of interrogational torture fulfils public-safety obligations that are more stringent than the officials' duties to eschew any use of such torture, their actions are seriously wrong. The officials may have acted correctly, but the correct course of conduct in a predicament of moral conflict is a breach of a moral duty. Given the stringency of the torture-eschewing obligation, the withholding of legal sanctions in the aftermath of a breach of that obligation would be morally untenable.²⁷ In other words, in combination with my conception of the moral role of law in a liberal-democratic system of governance, my emphasis on moral conflicts impels me quite smoothly toward my position on the appropriate legal status of torture.

These closing paragraphs laconically anticipate matters that will be discussed at much greater length in the rest of this book. The abstract points of moral philosophy expounded in the present chapter will surface again and again in this book as we turn our attention to the complexities of torture and of torture's wrongness. While the abstract points inform and orient my subsequent chapters' discussions, they are also crucially reinforced by those discussions. What this book aims to show is that its strongly deontological outlook—with its attentiveness to moral conflicts—is vital for any effort to do justice philosophically and morally to the fiendish problem of torture.

²⁶ Moore—1997, 733–4—does support a blanket legal prohibition on the use of torture against innocents (such as the young children of terrorists). However, he regards as unfair the imposition of sanctions on any officials who have averted calamities by resorting to such torture, and he acquiesces in the punishment of those officials only because he thinks that a policy of legally approving their actions against innocents would be too likely to lead to many instances of unjustifiable torture.

²⁷ As will become clear in Chapter 5, I am not assuming that the appropriate sanctions will always be criminal, nor am I assuming that they will always be imposed on individuals; in some cases, the sanctions should be imposed solely on the collectivities on whose behalf the individuals have resorted to torture.