

---

MORAL  
KNOWLEDGE?

---

New Readings in Moral  
Epistemology

---

Walter Sinnott-Armstrong  
Mark Timmons

*New York* *Oxford*  
OXFORD UNIVERSITY PRESS  
1996

To Liz  
W. S-A.

To John, Bryan, and Andy  
M.T.

Oxford University Press

Oxford New York  
Athens Auckland Bangkok Bombay  
Calcutta Cape Town Dar es Salaam Delhi  
Florence Hong Kong Istanbul Karachi  
Kuala Lumpur Madras Madrid Melbourne  
Mexico City Nairobi Paris Singapore  
Taipei Tokyo Toronto

and associated companies in  
Berlin Ibadan

Copyright © 1996 by Oxford University Press, Inc.

Published by Oxford University Press, Inc.,  
198 Madison Avenue, New York, New York 10016

Oxford is a registered trademark of Oxford University Press

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, electronic, mechanical,  
photocopying, recording, or otherwise, without the prior  
permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data  
Moral knowledge? : new readings in moral  
epistemology / edited by Walter Sinnott-Armstrong, Mark Timmons.  
p. cm.

Includes bibliographical references.  
ISBN 0-19-508989-8 (pbk.)  
ISBN 0-19-508988-X (cloth)

1. Ethics. I. Sinnott-Armstrong, Walter, 1955-  
II. Timmons, Mark, 1951-  
BJ1012.M633.1996 170—dc20  
95-44287

Chapter 6, *Foundationalism and Coherentism in Ethics*,  
Copyright © 1996 by R.M. Hare

1 3 5 7 9 8 6 4 2

Printed in the United States of America  
on acid-free paper

## Preface

In moral philosophy classes, professors usually explain many different moral theories, and then their students ask, "But how can I know which theory to accept?" or "How can I be justified in believing one moral claim instead of another?" These persistent questions are addressed in the field of moral epistemology and in these pages. No single answer is given by this volume. Instead, the editors gathered contributors who represent a wide variety of perspectives. Without assuming any background beyond an introduction to ethics, each contributor discusses a general kind of answer to the questions of moral epistemology and then defends his or her own particular answer within that general kind. Together, these essays should serve both to introduce the field of moral epistemology and to stimulate and contribute to its growth and development.

A general introduction is provided by sections 1–4 of the first chapter, by Walter Sinnott-Armstrong. Sinnott-Armstrong goes on to develop his own account of justification and his own limited version of moral skepticism. Theories of knowledge and justification depend crucially on the notion of truth, which is discussed in the following chapters, by Peter Railton and Simon Blackburn, who argue for and against moral realism, respectively.

The remaining essays can be seen as various responses to the skeptical regress argument. This argument is described in more detail in Sinnott-Armstrong's chapter, but the basic problem is simple: If a moral belief cannot be justified without depending on an inference from other beliefs, and if those other beliefs must also be justified, then justification of a moral belief seems to require a vicious infinite regress of one justification after another. In response, Robert Audi defends a version of moral intuitionism that claims that moral beliefs can be justified independently of any inference. Geoffrey Sayre-McCord then defends a version of coherentism that makes justification depend on inferential relations throughout the believer's system of beliefs, including the very belief to be justified. Richard Hare next tries to reconcile coherentists with foundationalists (including intuitionists), and also discusses the importance of understanding moral language for justifying moral beliefs. Richard Brandt then attempts to show how moral claims can be justified by appealing to science, especially cognitive psychology. Christopher Morris follows by defending a contractarian justification that grounds morality in a nonmoral notion of individual rationality; and David Copp develops a related approach that instead refers to rationality in choices by societies. Margaret Walker

also emphasizes the social dimension of morality, although from a feminist perspective, and argues that a certain kind of negotiation within society can lead to justification for some moral claims. Finally, Mark Timmons defends the contextualist view that some moral beliefs can be justified by reference to unjustified moral assumptions in certain contexts. Although there are other possible moves, these essays cover most of the main approaches to moral epistemology. Further readings in these and other theories about moral epistemology can be found in the annotated bibliography prepared by Mitch Haney.

Despite the contributors' disagreements, this book is truly a collaborative effort. Early versions of each essay were mailed to each contributor. We then held a conference at Dartmouth College, where everyone had a chance to comment on everyone else's work. A flurry of written exchanges followed. We all benefitted from the close scrutiny of the other contributors, as well as of the other participants at the conference.

Many others have also contributed to this project. Generous funding for the Moral Epistemology Conference at Dartmouth College was provided by the National Endowment for the Humanities, the Mellon Foundation, the Eunice and Julian Cohen Professorship for the Study of Ethics and Human Values, and the Gramlich lecture fund administered by the Philosophy Department at Dartmouth College. Barbara Hillinger and Julie Wright of the Dartmouth Institute for Applied and Professional Ethics helped to organize the conference. Research assistance was provided by Malia Brink through Dartmouth's Presidential Scholar Program. We are extremely grateful to all of these supporters.

*Hanover, N.H.*  
*October 1995*

W.S.-A.  
M.T.

## Contents

- Contributors, ix
- 1 Moral Skepticism and Justification, 3  
*Walter Sinnott-Armstrong*
  - 2 Moral Realism: Prospects and Problems, 49  
*Peter Railton*
  - 3 Securing the Nots: Moral Epistemology for the Quasi-Realist, 82  
*Simon Blackburn*
  - 4 Intuitionism, Pluralism, and the Foundations of Ethics, 101  
*Robert Audi*
  - 5 Coherentist Epistemology and Moral Theory, 137  
*Geoffrey Sayre-McCord*
  - 6 Foundationalism and Coherentism in Ethics, 190  
*R. M. Hare*
  - 7 Science as a Basis for Moral Theory, 200  
*Richard B. Brandt*
  - 8 A Contractarian Account of Moral Justification, 215  
*Christopher W. Morris*
  - 9 Moral Knowledge in Society-Centered Moral Theory, 243  
*David Copp*
  - 10 Feminist Skepticism, Authority, and Transparency, 267  
*Margaret Urban Walker*
  - 11 Outline of a Contextualist Moral Epistemology, 293  
*Mark Timmons*
  - 12 An Annotated Bibliography on Moral Epistemology, 326  
*Mitchell R. Haney*
- Index, 338

## Moral Skepticism and Justification

*Walter Sinnott-Armstrong*

In 1993, Dr. Jack Kevorkian helped Ali Khalifi commit suicide. Many people believe that this act was morally wrong, but others believe just as strongly that it was not. More generally, many people believe that sodomy is immoral, but others do not. And Kantians and utilitarians disagree about whether it would be morally wrong to lie or break a promise to someone just to make someone else happier. Such moral disputes raise many questions. If I have not yet formed any opinion about the morality of capital punishment, for example, how should I decide what to think? If I come to believe that capital punishment is immoral, is this belief justified? Can other people be justified in believing the contrary? Can anyone know whether capital punishment is immoral? How?

Similar questions arise even when there is no disagreement. Almost everyone agrees that torturing babies just for fun is immoral. But is this common belief justified? Do we know that it is true? What could we say to someone who did disagree? Could we show them that their beliefs are false or unjustified? How?

Such issues lie at the heart of moral epistemology. Whereas substantive ethics is about what is morally right or wrong, moral epistemology asks whether and how anyone can know or be justified in believing that something is morally right or wrong. These questions lead into deep issues about the nature of morality, language, metaphysics, and justification and knowledge in general. All of these issues can be illuminated by studying their implications for moral epistemology.

But the issue of moral justification is not just theoretical. It also has practical importance. Debates about when, if ever, an employee health plan should pay for abortions often turn on disputes about whether someone can know that abortion is morally wrong. Also, many legal theorists argue that judges should not overturn laws on the basis of their own moral beliefs, because those moral beliefs should not be trusted when the majority disagrees. And

educational policies often depend on views about whether values can be justified well enough to be taught in the public schools of a free society. In these and other ways, our public debates and institutions are deeply affected by our views on whether moral beliefs can be justified.

Even in our personal lives, we often need to decide which moral claims to believe, how much confidence to place in them, and what to think of people who believe otherwise. Suppose a friend has a secret affair with a married man. If you believe that the affair is morally wrong, you have to decide what, if anything, to do about it. Your decision might hinge on whether you think that your moral belief is justified well enough.

In general, when a moral problem is serious, most people want to have some belief about it, and they do not want this moral belief to be arbitrary or unjustified. So they want to form justified moral beliefs. The question is whether they can get what they want. This is the basic question for moral epistemology.

### What Is Meta-Ethics?

In order to better understand the field of moral epistemology, it is useful to locate it within the larger territory of moral theory. Any division of moral theory is bound to be controversial, but some framework can help in comparing various views. For this purpose, moral theory is often divided, first, into substantive ethics and meta-ethics.

Substantive ethics or normative ethics includes nontautologous claims and beliefs about what is morally right or wrong, what is morally good or bad, what morally ought or ought not to be done, and so on. These claims and beliefs might be about acts, states of character, persons, policies, institutions, laws, and so on. They might be about particular cases or about general kinds; and they might be combined into systems, such as utilitarianism. Anyone who makes or implies any such claim is to that extent within substantive ethics.


Meta-ethics then asks about the nature or status of substantive moral claims and theories. One prominent area of meta-ethics has been the study of moral language. When a speaker says, "Abortion is immoral," one might ask what this sentence *means*, what *effect* this utterance causes or is intended to cause, or what *speech act* is performed. Such questions fall under *moral semantics* or, more broadly, *moral linguistics*. The focus on moral language has produced some disdain for meta-ethics, but there is much more to meta-ethics than moral linguistics. Meta-ethics also includes *moral ontology*, which asks about the metaphysical status of moral properties and facts, if any. There are also debates about how to *define* morality as opposed to self-interest, religion, law, custom, and so on.

Moral epistemology is yet another area of meta-ethics. Moral epistemologists study justification and knowledge of substantive moral claims and beliefs. Of course, one cannot determine whether a claim is justified or known to be true if one has no idea what that claim means, so moral epistemology depends

in some ways on moral semantics. It also depends on moral ontology and on the definition of morality. Nonetheless, moral epistemology differs from other areas of meta-ethics in that it focuses directly on justification and knowledge of morality and brings in other theories only insofar as they are relevant to these central concerns.

Moral epistemology is also distinct from substantive ethics insofar as theories in moral epistemology are supposed to be neutral among competing substantive moral views. When people disagree about the morality of abortion, they often want a method to resolve their dispute or a test of whether either view is justified. Moral epistemologists study and sometimes propose such methods and tests. In order to avoid begging the question, these methods and tests must be neutral among the views under dispute. Of course, some theories in moral epistemology are not really so neutral; and maybe no theory in moral epistemology can be neutral among all possible substantive moral views. Even so, some theories in moral epistemology can be neutral among the competitors in a particular dispute; and then they might help in choosing among those alternatives.

These divisions of moral theory remain questionable, but they provide a rough picture of some of the various issues that arise about morality:



### What Is Moral Epistemology?

Moral epistemology is simply epistemology applied to substantive moral claims and beliefs. Epistemology is the study of knowledge and justification in general. It asks whether, when, and how claims or beliefs can be justified or known or shown to be true. Moral epistemology then asks whether, when, and how substantive moral beliefs and claims can be justified or known or shown to be true.

To determine whether moral knowledge is possible, we need to determine what is necessary for knowledge. I don't know where President Clinton is right now if I do not have any belief about where he is, or if I believe that he is somewhere other than where he is, or if my belief is true but just a guess. Examples like these have led many philosophers to define knowledge as justified true belief. If this also holds for morality, one cannot know that stealing is immoral unless it is true that stealing is immoral, one believes

that stealing is immoral, and one's belief is justified. This tradition has been questioned because justified true belief does not seem sufficient for knowledge in some cases,<sup>1</sup> and some philosophers argue that justification is also not necessary for knowledge.<sup>2</sup>

But regardless of how knowledge is analyzed, justification remains a distinct and important issue for epistemology. In moral epistemology in particular, many theories study the conditions under which moral claims or beliefs are justified or under which a person is justified in believing a moral claim. Other theories focus on processes and procedures that are supposed to justify moral beliefs and claims, or to be epistemically responsible ways to form moral beliefs. Some philosophers also develop methods for showing that moral claims or beliefs are true or justified. All of these issues about justification are important, but this essay will focus on the questions of which conditions must be met for a person to be justified in believing a substantive moral claim; and when, if ever, those conditions are met.<sup>3</sup>

### What Is Moral Skepticism?

The main challenges in moral epistemology are posed by moral skeptics. However, many different views have been described as moral skepticism. The best way to explain the particular kind of moral skepticism that will be our topic is to contrast it with the other views that sometimes go by the same name.

Two kinds of moral skepticism are directly relevant to moral epistemology, so they can be seen as versions of epistemological moral skepticism. The first is the claim that nobody is ever justified in believing any substantive moral claim. This will be called *moral justification skepticism*. The second is the claim that nobody ever knows that any substantive moral claim is true. This can be called *moral knowledge skepticism*. If knowledge implies justification, as is often assumed, then moral justification skepticism implies moral knowledge skepticism. But even if knowledge requires justification, it does not require only justification, so moral knowledge skepticism does not imply moral justification skepticism.

Both of these views must be distinguished from the meta-ethical claim that no substantive moral belief or claim is either true or false. Some philosophers argue that claims like "Cheating is morally wrong" are neither true nor false, because they resemble expressions of emotion (such as "Boo Nicks") or prescriptions for action (such as "Go Celtics"), which are neither true nor false. Such theories rest on linguistic analyses, so I will call them *linguistic moral skepticism*, although a more common name is *noncognitivism*.<sup>4</sup>

Such analyses of moral language face serious problems. Critics argue that moral judgments do not simply express emotions or prescribe acts but also or instead make claims about the moral properties of those actions. When Ronald Reagan said, "Abortion is morally wrong," he wanted to say something true about abortion. He did not want only to express personal emotions or only to tell people how to act. Moreover, it would be difficult to explain the

role of moral claims in logical inferences if moral claims were analyzed solely as expressions or prescriptions.<sup>5</sup> In response to such arguments, most recent defenders of emotivism and prescriptivism grant that moral beliefs can be true or false at least in some minimal way, possibly captured by a semantic or redundancy theory of truth, so they are not linguistic moral skeptics on my definition.<sup>6</sup> The issue of truth in morality deserves much more careful attention, but I will assume henceforth that moral claims and beliefs can be true or false at least in a minimal way.

Yet another kind of moral skepticism still might be defensible. Error theorists argue that moral claims try to say something about moral properties (such as moral wrongness) of actions, but there are no properties of the kind that is claimed, so all positive moral claims are false.<sup>7</sup> Because ontology is the study of what exists, the view that no moral properties or facts exist can be called *ontological moral skepticism*.

Despite their differences, linguistic moral skeptics and error theorists agree that no moral claim or belief is true. This has direct implications for moral knowledge skepticism. If knowledge implies truth, as is usually assumed, and if moral claims are never true, then there is no knowledge of what is moral or immoral. However, the converse implication fails: even if some moral claims are true, it still might not be possible for anyone to know that a moral claim is true. One reason might be that moral claims cannot be justified, and knowledge requires justification in addition to truth. So, even if linguistic and ontological moral skepticism are false, moral justification skepticism and moral knowledge skepticism still might be true.

Yet another view is often described as moral skepticism. It is common to ask, "Why should I be moral?" Most people admit that there is *sometimes* a reason to be moral, but the controversial claim is that there is *always* a reason to be moral (or to do what is morally required).<sup>8</sup> A denial of this claim is often seen as a kind of moral skepticism. I will call it *practical moral skepticism*.

This view resembles moral justification skepticism in that both skepticisms deny a role to reasons in morality. However, moral justification skepticism is about reasons for *belief*, whereas practical moral skepticism is about reasons for *action*. Moreover, practical moral skepticism denies that there is *always* a reason for moral action, whereas moral justification skepticism denies that there is *ever* an adequate reason for moral belief. Because of these differences, practical moral skepticism does not imply moral justification skepticism. If we are ever justified in believing that an act is immoral, then moral justification skepticism is false. But even if so, and even if there is a reason to act morally in that case, practical moral skepticism still might be true, because there might not be any reason to act morally in other cases.

In sum, then, we need to distinguish the following kinds of moral skepticism:

*Linguistic moral skepticism* = no moral belief or claim is true or false.

*Ontological moral skepticism* = no moral facts or properties exist.

*Practical moral skepticism* = there is not always a reason to be moral.

*Moral knowledge skepticism* = no moral belief or claim is known to be true.  
*Moral justification skepticism* = no moral belief or claim is justified.

The last view will be the main topic from now on, and I will refer to it simply as moral skepticism.

Philosophers who accept this moral skepticism need not be any less motivated to be moral, nor need they have any less reason to be moral, than their opponents. Moral skeptics can hold moral beliefs just as strongly as nonskeptics. Moral skeptics can even believe that their moral beliefs are true by virtue of corresponding to an independent reality. All that moral skeptics must deny is that their (or anyone's) moral beliefs are justified in the relevant way, but this is enough to make moral skepticism very controversial and important.

### Is Moral Skepticism True?

There is an initial presumption against moral skepticism. Most people hold many moral beliefs and think that they are justified. It seems obvious that it is morally wrong to hit your opponent just to get her back for beating you in a game.<sup>9</sup> Furthermore, people often say that they *know* that something is morally wrong. If a father punishes his child by spanking her severely, his neighbors might say, "I know it is wrong for him to spank her, but I don't know what I should do about it." We also often say that people *ought* to know what is immoral. If a stock broker is caught passing inside information but claims that he did not know that this was unethical, many people might respond, "He should have known better," where this does not mean just "He should have known that he would be caught" but rather "He should have known that what he did was unethical." But he could not have known this, according to moral knowledge skeptics; therefore, this moral skepticism conflicts with these common ways of talking and thinking. None of these examples is conclusive, and common beliefs might be wrong on these matters. Yet such considerations create a presumption against moral skepticism. In order to override this presumption, moral skeptics need at least some argument.<sup>10</sup>

Many different kinds of argument have been given for moral skepticism. Some moral skeptics argue that no moral claim or belief can be true (linguistic or ontological moral skepticism), and they infer that no moral claim or belief can be known to be true (moral knowledge skepticism). Others argue that people disagree and make so many mistakes about morality that our moral belief-forming processes are not reliable enough to result in justified beliefs or knowledge. Still others argue that nobody can know what is morally right or wrong unless they are fully informed and impartial, but nobody can be fully informed and impartial about moral issues. And some argue that we cannot know that any claim is true unless the truth of that claim is necessary for the best explanation of some independent fact, but moral truths are never necessary for the best explanation of any nonmoral fact.<sup>11</sup>

Each of these arguments deserves serious consideration, but here I will feature two other arguments for moral skepticism. These two arguments are versions of well-known arguments for general skepticism, which is the view that no claims or beliefs of any kind can be justified or known. I will apply these arguments to morality in particular. Although these arguments can also be used outside morality, they cannot be dismissed just because they are general. If the problems raised by these arguments cannot be solved at least in morality, then we cannot be justified in believing any moral claims, so moral skepticism is true.

### The Regress Argument

The first featured argument is a regress argument, which derives from Sextus Empiricus.<sup>12</sup> Its goal is to lay out all of the ways in which a person might be justified in believing something, and then to argue that none of them works.

The regress argument begins by distinguishing two ways to be justified:

A person *S* is *inferentially* justified in believing a moral claim that *p* if and only if what makes *S* justified is at least in part that *S* believes something from which *S* does or could infer *p*.

Here *S*'s being justified depends on *S*'s ability to infer *p* from a belief of *S*, so *S* would not be justified if *S* could not draw such an inference. In contrast,

A person *S* is *noninferentially* justified in believing a moral claim that *p* if and only if *S* is justified but not inferentially justified in believing that *p*.

Here *S*'s being justified does not depend on any ability of *S* to infer *p* from a belief of *S*. What makes *S* justified then might be something that is independent of any actual or available belief or inference, such as a nondoxastic perceptual or emotional state; or it might be nothing at all. In this last case, *S* is *directly* justified.

The first premise of the regress argument then claims that these are the only ways to be justified:

- (1) If any person *S* is justified in believing any moral claim that *p*, then *S* must be justified either inferentially or noninferentially.

The moral skeptic goes on to deny both of these possibilities. First,

- (2) No person *S* is ever noninferentially justified in believing any moral claim that *p*.

In other words, *S* must be able to infer *p* from some other beliefs of *S* in order to be justified in believing that *p*.

Premise (2) is often supported by examples. Thelma believes that it is morally wrong to eat meat; Louise asks "Why?" If Thelma has no reason at all for her moral belief, then she does not seem justified in believing it. Thelma's belief seems to be the kind that needs to be justified by inferring it from some reason. If there is nothing special about this moral belief, then the same point applies to every moral belief. Moral skeptics conclude that noninferential justification never works, as premise (2) claims.

Premise (2) is denied by moral intuitionists, who claim that some moral beliefs are justified independent of any ability to infer them from any other beliefs.<sup>13</sup> What is believed is then sometimes called self-evident or evident by itself. Moral skeptics see this as dogmatism, and dogmatism can be used on both sides of any issue. If Thelma could be noninferentially justified in believing that eating meat is wrong, then Louise could also be noninferentially justified in believing that eating meat is *not* wrong, even if neither can infer her belief from any reason.<sup>14</sup> Moreover, Nick could be noninferentially justified in believing that it is morally wrong to eat artichokes or to step on cracks in sidewalks. Moral skeptics claim that there is no way to break such impasses or to avoid such absurdities if any noninferential justification is admitted.

Given (1) and (2), the moral skeptic can draw an intermediate conclusion:

- (3) If any person *S* is justified in believing any moral claim that *p*, then *S* must be justified inferentially.

This means that *S* must be able to infer *p* from some other beliefs held by *S*. But which other beliefs? There are two main possibilities:

- (4) If any person *S* is inferentially justified in believing any moral claim that *p*, then *S* must be justified either by an inference with some moral premises or by an inference without any moral premises.

Inferences can contain a mixture of moral and nonmoral premises, but the question here is whether *S* can be justified by an inference from nonmoral premises alone without any moral premise at all.

An inference from nonmoral premises might seem to have advantages if nonmoral premises are easier to justify than moral premises. However, even if the nonmoral premises are secure, it is still not clear how nonmoral premises alone could be enough to justify moral conclusions. Some inferences of this kind start from premises that are supposed to be descriptive rather than normative.<sup>15</sup> In a simple example, premises like

- (ii) The act has feature *F* (such as that it breaks a promise);  
 (iii) The context is of kind *C* (such as where people depend on each other);

- (iv) People in group *G* would have reaction *R* (such as disapproval)<sup>16</sup>

might be used to justify a conclusion like

- (v) The act is morally wrong.

Moral skeptics claim that all such inferences are incomplete because they depend on a bridge principle like

- (i) Any act is morally wrong if it has feature *F* in context *C* and causes reaction *R* in group *G*.

The point is not only that the argument from (ii)–(iv) to (v) is not formally valid without (i), but also that (i) is needed because (ii)–(iv) would not justify (v) if no bridge principle like (i) were true. Yet bridge principles like (i) seem to be substantive moral claims, so they need to be justified themselves. Moral skeptics conclude that nonnormative premises like (ii)–(iv) are never enough by themselves to justify any moral beliefs.

Some philosophers have proposed much more complex arguments to justify moral beliefs by inferring them from other kinds of nonnormative premises. These include premises about metaphysics<sup>17</sup> and various sciences, such as sociobiology,<sup>18</sup> developmental psychology,<sup>19</sup> and the psychology of motivation.<sup>20</sup> However, moral skeptics respond that these more complex arguments still depend on bridge principles that are not morally or normatively neutral, and these bridge principles need to be justified themselves. If so, then no moral belief can be justified by inferring it from purely nonnormative premises.

A related attempt to justify moral beliefs uses premises that are not moral but still are normative in other ways. This approach is taken by many contractarians and others.<sup>21</sup> Contractarians usually start with nonmoral normative premises about who is rational (and sometimes also about who is impartial), argue that these people under certain circumstances would agree to certain moral rules (or norms or standards), and conclude that the corresponding moral beliefs are true (or justified). Even if these premises can be established—which is not easy—such arguments still need another bridge principle to relate the nonmoral premises to the moral conclusion. In one case, the needed bridge principle claims that an act is morally wrong if it violates a rule that would be accepted by all rational, impartial people in appropriate circumstances. This bridge principle might seem innocuous at first, but that is because it has little content until "rational," "impartial," and "appropriate circumstances" are defined. Once these terms are defined, the bridge principle becomes questionable. The problem is that different acts come out morally wrong when this contractarian framework is filled in with different background theories of rationality, impartiality, and appropriate circumstances. This flexibility puts contractarians in a dilemma. If the background theories are *not* detailed enough to yield specific moral beliefs, then the contractarian form



of argument will not justify those moral beliefs. On the other hand, if the background theories *are* detailed enough to yield specific moral beliefs, then anyone who wants to deny those moral beliefs can (with equal justification) deny those aspects of the background theories that are needed to yield those moral beliefs. Either way, the contractarian argument will fail to justify the moral belief, according to moral skeptics. There are other arguments from the same or other nonmoral norms, but they seem to run into similar problems.

If nonnormative premises cannot justify moral beliefs, and nonmoral normative premises also cannot justify moral beliefs, then

- (5) No person *S* is ever justified in believing any moral claim that *p* by an inference without any moral premises.

When (5) is added to the previous premises (3)–(4), the moral skeptic can draw another intermediate conclusion:

- (6) If any person *S* is justified in believing any moral claim that *p*, then *S* must be justified by an inference with some moral premise.

In short, moral beliefs must be justified by moral beliefs.

This creates a problem. Although the justifying beliefs must include some moral beliefs, not just any moral beliefs will do:

- (7) No person *S* is ever justified in believing a moral claim that *p* by an inference with a moral premise unless *S* is also justified in believing that moral premise itself.

Premise (7) is denied by some contextualists, who claim that even if a moral belief is not itself justified, if it is shared within a certain social context, then it can be used to justify other moral beliefs.<sup>22</sup> However, the fact that a moral belief happens to be held by everyone in a group is not sufficient to make anyone justified in believing it or in believing anything that follows from it. To see this, suppose that everyone in a small town believes that it is immoral for a black person to marry a white person. They infer that it is immoral for Ray, who is black, to marry Terry, who is white. This conclusion is not justified if the general principle is not justified. Indeed, every belief would be justified if a belief could be justified simply by inferring it from an unjustified belief, since every belief can be validly inferred from itself. The only way to avoid such absurd results, according to moral skeptics, is to accept (7).

But then how can the moral premises be justified? Given (6)–(7), the moral premises must be justified by inferring them from still other moral beliefs that must also be justified by inferring them from still other moral beliefs, and so on. To justify a moral belief thus requires a branching tree or chain of justifying beliefs or premises. This justifying chain can take only two forms, so that

- (8) If any person *S* is justified in believing any moral claim that *p*, then *S* must be justified by a chain of inferences that either goes on infinitely or includes *p* itself as an essential premise.

The latter kind of chain is usually described as circular, although its structure is more complex than a simple circle.

Moral skeptics then deny that either kind of chain can justify any moral belief. First,

- (9) No person *S* is ever justified in believing any moral claim that *p* by a chain of inferences that includes *p* as an essential premise.

Any argument that includes its conclusion as a premise will be valid and will remain valid if other premises are added. However, anyone who doubts the conclusion will have just as much reason to doubt the premise. So skeptics claim that nothing is gained when a premise just restates the belief to be justified.

Premise (9) is opposed by moral coherentists. Recent coherentists have emphasized that they are not inferring a belief from itself in a linear way. Instead, a moral belief is supposed to be justified because it coheres in some way with a body of beliefs that is coherent in some way.<sup>23</sup> Still, moral skeptics deny that coherence is enough for justification. One reason is that the internal coherence of a set of beliefs is not evidence of any relation to anything outside the beliefs. Moreover, every belief—no matter how ridiculous—can cohere with some body of beliefs that is internally coherent. If Bill argues that marriage between people with different skin colors is immoral, and Claire points out that the same reasons would show that marriage between people with different hair colors is immoral, then most people would see this as a reason to reject Bill's original belief. Nonetheless, Bill could hold on to his original belief that people with different skin colors should not be married if he also accepts that people with different hair colors also should not be married. He could even derive both beliefs from a general principle against mixing colors. Bill's resulting set of beliefs could be just as coherent internally as it would be if Bill saw neither kind of marriage as immoral. But Bill still does not seem justified in believing that both kinds of marriage are immoral. So coherence does not seem to be enough to make moral beliefs justified.

The final possible form of justification is an infinite chain. Moral skeptics, of course, claim that

- (10) No person *S* is ever justified in believing any moral claim that *p* by a chain of inferences that goes on infinitely.

Someone who denied this premise could be called a moral infinist, but nobody clearly endorses this approach.<sup>24</sup>

Now the moral skeptic can draw a final conclusion. Premises (8)–(10) imply that

- (11) No person is ever justified in believing any moral claim,

which is moral skepticism. This conclusion might seem implausible, but the regress argument is valid, so its conclusion can be escaped only by denying one of its premises.

Of course, every step along the way could be opposed. Moral intuitionists deny premise (2). Moral naturalists and contractarians deny premise (5). Moral contextualists deny premise (7). Moral coherentists deny premise (9). And a moral infinitist would deny premise (10). The regress argument thus provides a useful way of classifying theories in moral epistemology, regardless of whether it succeeds in establishing moral skepticism.

The most important question, however, is whether the regress argument shows that moral skepticism is true. I have tried to say enough to give some initial plausibility to the premises of the regress argument, but each premise deserves much more careful attention. In particular, because they are about how moral beliefs can be justified, these premises and their denials cannot be assessed properly without first determining what it means for a moral belief to be justified. I will address that issue, but first I want to look at the other main argument for moral skepticism. Then I will distinguish different ways for a belief to be justified. Finally, I will use my distinctions to evaluate both arguments for moral skepticism.

### *Skeptical Hypothesis Arguments*

The second featured argument, which derives from Descartes,<sup>25</sup> starts from the common experience of being deceived. For example, yesterday I was driving down a strange road, and I thought I saw a lake in the distance, but my wife said it was just a mirage. Always curious, we turned toward it and found that I was right. We even took a swim. But what if we had not checked it out? Then I would not be justified in believing that what I saw was a lake, because I could not rule out the contrary hypothesis that it was a mirage. My belief would also not be justified if I could not rule out the possibility that it was a river rather than a lake. The general principle seems to be that I am not justified in believing something if I cannot rule out all contrary hypotheses.<sup>26</sup>

Skeptics then introduce hypotheses that cannot ever be ruled out. The famous Cartesian hypothesis is of a demon who deceives me in all of my beliefs about the external world, while also ensuring that my beliefs are completely coherent. If there is such a deceiving demon, then there really is no lake when I think I see one. Nobody claims that such a deceiving demon actually exists, but it is possible.<sup>27</sup> And this possibility cannot be ruled out by any beliefs or sense experiences. Even when I seem to be swimming in a lake, my sensations might all be caused by a deceiving demon. Since there is no way to rule out this skeptical hypothesis, my beliefs about the lake are not justified, according

to the above principle. And there is nothing special about my beliefs about the lake. Everything I believe about the external world is incompatible with the deceiving demon hypothesis. Skeptics conclude that no such belief is justified.

This kind of argument can be applied to moral beliefs in several ways. First, another demon might deceive us about morality (either alone or in addition to other topics). Most people believe that it is morally wrong to torture babies just for fun, but we might be deceived in our beliefs that babies feel pain or that they have moral rights. A demon might make us believe that some creatures have moral rights, when really they do not, although other things do; or a demon might make us believe that some creatures have moral rights when really nothing does (just as a demon might deceive us into believing that some women are witches when really none are). Nobody actually claims that a demon distorts our moral beliefs in these ways, but such deception still seems possible. And such a deceiving demon cannot be ruled out by the fact that it seems obvious that babies have moral rights, any more than Descartes's deceiving demon could be ruled out by the fact that it seems obvious to me when I am swimming in a lake.

Moral skeptics do not have to depend on deceiving demons. Different arguments with the same structure can be constructed with skeptical hypotheses that are peculiar to morality. Consider:

Moral Nihilism = Nothing is morally wrong.<sup>28</sup>

Moral nihilism seems consistent in itself as well as logically compatible with all nonmoral facts and with their best explanations.<sup>29</sup> Just as it would beg the question to use common beliefs about the external world to rule out a deceiving demon hypothesis, so it would also beg the question to argue against moral nihilism on the basis of common moral beliefs—no matter how obvious. Thus, neither logic nor nonmoral beliefs nor moral beliefs can rule out moral nihilism. However, moral nihilism implies that it is not morally wrong to torture babies just for fun. So, according to moral skeptics, one must be able to rule out moral nihilism in order to be justified in believing that torturing babies just for fun is morally wrong. Moral skeptics conclude that this moral belief is not justified. More precisely:

- (1') I am not justified in believing that moral nihilism is not true.
- (2') I am justified in believing that "It is morally wrong to torture babies just for fun" (*p*) entails "Moral nihilism is not true" (*q*).
- (3') If I am justified in believing that *p*, and I am justified in believing that *p* entails *q*,<sup>30</sup> then I am justified in believing that *q*.
- (4') Therefore, I am not justified in believing that it is morally wrong to torture babies just for fun.

This moral belief is not especially problematic in any way. Indeed, it seems as obvious as any moral belief. So the argument can be generalized to cover any moral belief. Moral skeptics conclude that no moral belief is justified.<sup>31</sup>

There are two main responses to such skeptical hypothesis arguments. First, some antiskeptics deny (1') and claim that skeptical hypotheses can be ruled out as meaningless, because the meaning of the phrase "morally wrong" depends on judgments about paradigm cases, such as "It is morally wrong to torture just for fun."<sup>32</sup> However, when a claim really is meaningless, such as "The square root of pine is tree," then its denial is meaningless, too. But the denial of "Torturing for fun is not morally wrong" is "Torturing for fun is morally wrong." The latter claim is not meaningless, so neither is the former. Antiskeptics might reply that "Torturing for fun is morally wrong" is more like "Cows are animals," whose denial is not meaningless but is incoherent, because the meaning of "animal!" can be given as "Something like a cow or a fish or . . ." But then, if there is any question about whether something is an animal, there is just as much question about whether it is a cow. Similarly, if moral wrongness is defined as "Something like torturing just for fun or . . .", then, if there is any question about whether an act is morally wrong, there is just as much question about whether to count it as torturing just for fun in the sense that is supposed to define "morally wrong." Moreover, this account cannot show that any morally wrong act exists. Compare someone who explains what a witch is by saying "Someone like the Wicked Witch of the West or . . ." This cannot show that any witch exists, so the analogous argument in morality also cannot show that there exists any act that is morally wrong. Finally, even if this conclusion did follow, it would be unclear why moral skeptics or nihilists must accept that moral wrongness can be defined as "Something like torturing just for fun or . . .", or that any claim like "Torturing for fun is morally wrong" is part of the common meaning of "morally wrong." So this approach cannot rule out moral nihilism.

Another common response is to deny premise (3'); which is a principle of *closure*. Since a belief entails the denial of every contrary hypothesis, this closure principle in effect says that I cannot be justified in believing *p* unless I am justified in denying every hypothesis contrary to *p*—that is, unless I can rule out *all* contrary hypotheses. This principle has been denied by relevant alternative theorists, who claim instead that only relevant hypotheses need to be ruled out. On this theory, if skeptical hypotheses are not relevant, then a belief that it is morally wrong to torture babies just for fun can be justified, even if the believer cannot rule out moral nihilism or deceiving demons.

But why would these skeptical hypotheses be irrelevant? They are consistent and contrary to the moral belief to be justified. One might argue that these hypotheses are irrelevant just because they cannot in principle be ruled out, and a moral belief is justified when the believer has done everything he or she can do to check it out. But that just begs the question against skepticism. Besides, hypotheses that cannot be ruled out often prevent beliefs from being justified. Consider Hannah, who is trying to figure out what is inside a wrapped birthday present. She knows that it is not a bike, because it is too small. She knows that it is not clothes, because of how it rattles when shaken. But it still might be either a puzzle or a game. If she believes that it is a game, her belief is not justified unless she can somehow rule out the hypothesis of a puzzle.

Now suppose that the evil donor put the present in an impenetrable box tied with an unbreakable ribbon. Hannah can't do anything else to find out what is in the box, but that does not make her justified in believing that the box contains a game. To the contrary, it keeps her belief from being justified. Similarly, if one cannot rule out the hypotheses of a deceiving demon and of moral nihilism, the fact that one does everything that one can do to determine whether it is morally wrong to torture babies just for fun does not make this moral belief justified, according to moral skeptics.<sup>33</sup>

Where does this leave us? The arguments for moral skepticism seem strong, and yet most people still reject their conclusion. Most people still think that they are justified in believing some moral claims, such as that it is morally wrong to torture babies just for fun. Moreover, such beliefs do seem to have a lot going for them: They seem obvious to many people, and they also connect well with other beliefs. So they seem justified to some extent. Nonetheless, this does not refute or even undermine the arguments for moral skepticism. Arguments cannot be refuted or undermined simply by denying their conclusions. Besides, there seems to be something right in the skeptics' arguments insofar as their premises, such as (1)–(10) and (1')–(3'), rest on general principles that seem plausible in many examples. So both sides seem to contain some truth. To see how this is possible, and to assess the arguments for moral skepticism, we need to look more carefully at what it means to call a moral belief justified.

### Kinds of Justification

When someone describes a moral belief and asks whether it is justified, many people often find themselves wanting to say both "Yes" and "No," even when all other facts are settled. This reaction is a signal that separate kinds of justification need to be distinguished.

One initial distinction is between negative or permissive justification and positive or supportive justification. Some people use "justified" in a very weak way. When they say that someone is justified in believing something, all they mean is that the believer does not violate certain standards, so she is epistemically permitted to believe it. This permissive use of "justified" occurs, for instance, when religious believers want to claim only that they are not irresponsible and do not have to give up their religious beliefs. But permissive justification cannot be what moral skeptics and their opponents disagree about. Moral skeptics need not claim that people should stop believing that it is morally wrong to torture babies just for fun. And skeptics' opponents claim more than that there is nothing wrong with moral beliefs. They claim that there is something right about those beliefs—some positive support for those beliefs. So debates about moral justification skepticism are about a more positive kind of justification.<sup>34</sup>

There are, of course, different degrees of positive support. Imagine, for example, that a murder was committed in the library, and the only people in

the building were Colonel Mustard, Professor Plum, and Miss Scarlet. A detective finds size thirteen boot prints in the mud outside the library window, and only Colonel Mustard wears size thirteen boots. This is *some* reason to believe that the Colonel committed the murder. Yet if there is no other evidence, the detective is not justified in believing that Colonel Mustard is guilty. The detective ought to withhold belief until stronger evidence is found. So a belief is not justified in the relevant way unless it has *enough* positive support that it ought to be believed.<sup>35</sup>

Thus, the general kind of justification that will be our topic is adequate positive support.<sup>36</sup> Within this general kind, we still need to draw two more distinctions among different kinds of positive support for beliefs. These two distinctions will be central to the rest of this paper.

### *Instrumental versus Evidential Justification*

The first distinction is between instrumental justification and evidential justification. It can be illustrated by a simple example.<sup>37</sup> Suppose that a benefactor will give me a million dollars if I believe that there are aardvarks on Mars, and I have a drug that will make me believe this. I am justified in taking the drug, so I am justified in making myself believe that there are aardvarks on Mars, so I am justified in being in this belief state, and my belief seems justified at least derivatively. However, I have no evidence for the content of my belief—that is, for what I believe. In cases like this, I will say that the believer has instrumental justification but not evidential justification.

The crucial difference is that evidential justification is tied to truth in a way that instrumental justification is not. The tie between evidence and truth is not simple, because one can be evidentially justified in believing a falsehood. For example, fresh aardvark droppings on Mars could give you evidence that there are aardvarks on Mars, even if the droppings had been transported from Earth without you knowing it. You also have evidence of aardvarks when you see an animal that looks just like an aardvark, even if it really is a disguised armadillo that fools you. So evidence need not guarantee truth. Nonetheless, evidence for a belief must be such that the content of the belief is more likely to be true given the evidence than without the evidence. So evidential justification is still tied to truth at least probabilistically.<sup>38</sup> In fact, evidential justification could be defined as adequate positive support that is tied to truth.<sup>39</sup>

In contrast, instrumental justification depends only on the beneficial *effects* of the mental *state* of belief.<sup>40</sup> It is being in the state of believing in aardvarks on Mars that gets me the million dollars. The acts of taking drugs that put me in that mental state are also instrumentally justified because those acts also have beneficial effects; but those acts have good effects only because they put me in the belief state; so the instrumental justification attaches to the belief state as well. It is not only the belief state but also its content that matters, because the effects of the belief state depend on its content. I do not

get the million dollars unless I have a belief with the right content. Nonetheless, it does *not* matter whether the content of the belief is true. My belief is instrumentally justified whether or not there really are aardvarks on Mars. It doesn't even matter whether this is probable. After all, it is just as easy to imagine that my benefactor will give me a million dollars for believing the opposite, or for believing anything else. In general, instrumental justification depends on the effects of a belief state, and the effects of a belief state do not affect the truth or probability of the belief's content,<sup>41</sup> so instrumental justification is not tied to truth even probabilistically.<sup>42</sup>

This same distinction applies as well to moral beliefs. If a benefactor will give me a million dollars to believe that nuclear deterrence is morally wrong, and if I have a drug that will make me believe this, then I might be instrumentally justified in getting and being in the mental state of believing this. Nonetheless, this instrumental justification still does not show that the content of this belief is true or even likely to be true, any more than in the nonmoral case. So it takes something else for this belief to be evidentially justified. What more it takes remains to be seen, but the million-dollar bribe is enough to show that instrumental justification is different from evidential justification even in ethics.<sup>43</sup>

### *Everyday versus Philosophical Justification*

Within each of these kinds, another distinction needs to be drawn between everyday and philosophical justification. This distinction depends on contrast classes, so that is the place to begin. Consider jumbo shrimp. Are jumbo shrimp large? An answer of "yes" or "no" would be too simple. Jumbo shrimp are large for shrimp, but they are not large for seafood. Whether one sees jumbo shrimp as large or not depends on whether one contrasts jumbo shrimp with other shrimp or with other seafood, such as salmon. Speakers need not have any very specific contrast class in mind, but it is silly to argue about whether jumbo shrimp are large if the arguers have in mind very different contrast classes.

Such contrast classes have been useful in many areas of philosophy, but most relevant here is their use in epistemology.<sup>44</sup> If epistemic justification is relative to contrast classes, then a belief can be justified relative to one contrast class even if it is not justified relative to another contrast class. This provides a new perspective on skepticism. Because I cannot rule out Descartes's deceiving demon, I cannot be justified in believing that I am sitting in my office as opposed to being deceived by a demon. But some claim that I can still be justified in believing that I am sitting in my office instead of standing in someone else's office. If so, the problem of general skepticism is reduced somewhat by analyzing justification in terms of contrast classes.

This approach can be extended to moral beliefs. Moral philosophers have produced many moral systems, but students often start with

(A) Act Utilitarianism = an act is morally wrong if and only if it fails to maximize the good for all,

and

(K) Kantian Absolutism = it is morally wrong to treat any person as a means only, even if doing so maximizes the good for all.<sup>45</sup>

The professor then describes some counterexample to act utilitarianism, such as cutting up one healthy innocent person and distributing his organs to save the lives of five unrelated people. If this act is morally wrong, and if Kantian absolutism implies that this act is morally wrong, and if act utilitarianism implies otherwise, then students who are justified in believing these premises seem to be justified in believing Kantian absolutism as opposed to act utilitarianism. However, even if this argument does accomplish this much, it would not justify anyone in believing in Kantian absolutism as opposed to some other system, such as

(R) Rule Utilitarianism = an act is morally wrong if and only if it violates a general rule such that everyone's obeying or accepting that rule would maximize the good for all,

if this also explains why the act in the example is morally wrong. More technically, one might be justified in believing moral system K out of the contrast class {A, K} without being justified in believing moral system K out of the larger contrast class {A, K, R}. Of course, another argument might justify one in believing K instead of R, and there are many more theories and many more arguments. Yet the fact that such arguments provide justification relative to some but not all alternatives is one reason to think that justification is relative to a contrast class.

The crucial question is: How far can such arguments go? Can they rule out *all* alternatives? It seems not. Recall the skeptical hypothesis of

(N) Moral Nihilism = Nothing is morally wrong.

This hypothesis is incompatible with act utilitarianism, rule utilitarianism, and Kantian absolutism. However, this hypothesis cannot be ruled out by semantics or by nonmoral facts. It also cannot be ruled out by common moral beliefs because any argument based on common moral beliefs begs the question against moral nihilism.

Must one be able to rule out moral nihilism to be justified in believing, say, rule utilitarianism? That is like asking whether jumbo shrimp are large. No simple answer is possible. One need not be able to rule out moral nihilism to be justified in believing rule utilitarianism as opposed to act utilitarianism or Kantian absolutism. However, one must be able to rule out moral nihilism to be justified in believing rule utilitarianism as opposed to moral ni-

hilism. Thus, if one cannot rule out moral nihilism, then one can be justified in believing R out of the contrast class {A, K, R} but not out of the larger contrast class {A, K, R, N}.

This relativization applies not only to general moral theories but also to specific moral beliefs. Suppose that a patient tests HIV +, but his doctor wants to lie to him about the test results, because learning the results would depress the patient. However, the hospital ethics committee discusses the case and concludes that it would be morally wrong for the doctor to lie and not to reveal the test results to the patient. This belief might be justified in contrast with many moral beliefs, such as that lying is usually wrong but not in this case, or that it would be morally wrong to lie in this case but not for the doctor to avoid lying by refusing to tell the patient anything about the test results. Nonetheless, if the ethics committee cannot rule out moral nihilism, then the committee's moral belief is not justified in contrast with moral nihilism. So this moral belief is also justified relative to some contrast classes but not others.

An impatient critic might ask: But is this belief *just plain justified*? That depends, of course, on what "justified" means when it is not qualified. This term might then refer to justification relative to an *unlimited* contrast class that includes all contraries of the belief said to be justified. If so, then anyone who cannot rule out moral nihilism also cannot be justified in holding any moral belief. However, when someone calls a belief justified and does not explicitly specify any contrast class, the speaker usually wants to claim only that this belief is justified out of the *relevant* contrast class. If so, then one cannot determine whether a belief is justified without qualification until one determines which contrast class is relevant.

Which contrast class is *seen* as relevant varies with context and purpose. This variation does not cause trouble, even when no contrast class is specified explicitly, as long as speakers assume contrast classes that are the same or similar enough. However, confusion results when a crucial alternative seems relevant in one context but not in another. Recall the doctor who asks the hospital ethics committee whether it is morally wrong to lie to his or her patient about the unfavorable test results. It would be very unusual for any member of the ethics committee to respond that morality is just an illusion, so nothing is morally wrong. If such nihilism were expressed, it would be dismissed quickly with disdain and without argument. Contrast this with a philosophy course that has studied moral nihilism, egoism, and skepticism all term. Many students have defended these positions in discussion. In the final paper, one student argues that it is morally wrong to lie, but does not even mention moral nihilism, egoism, or skepticism. This student would and should receive a low grade. Confusion arises if someone asks whether the committee's moral belief and the student's moral belief are justified. It seems to many people that the ethics committee is justified in concluding that the doctor should not lie, even though the committee never even considers moral nihilism. But the philosophy student does not seem justified in reaching the same conclusion because that student has no response to the same hypothesis. This

is paradoxical: if neither the philosophy student nor the doctor can rule out moral nihilism, how can this inability show that the philosophy student's belief is *not* justified when the committee's belief is justified?

The answer is that different contrast classes are assumed to be relevant in the different contexts. Hospital ethics committees usually assume that moral nihilism is not a relevant alternative. That is why they do not consider it. But philosophy classes often assume that moral nihilism is relevant. That is why they discuss it. Consequently, it is fine to say that the hospital ethics committee is justified in reaching its conclusion, but the philosophy student is not justified in reaching the same conclusion, even if neither can rule out moral nihilism. These two statements do not really conflict, because they assume different contrast classes to be relevant.<sup>46</sup> The paradox is thereby resolved. There appears to be a paradox only when different contrast classes are seen as relevant, but these different classes are not specified explicitly. This ability to resolve such paradoxes is another reason to relativize justification to contrast classes.

The point is not that the ethics committee is justified because it has more information or expertise or logical acumen than the philosophy student. Whether one is justified in believing something does depend on one's other beliefs and on one's intellectual abilities. As one gathers evidence and draws inferences, one comes to be justified in more beliefs. However, that kind of variation does not explain why the hospital ethics committee is justified when the philosophy student is not. In my example, the ethics committee and the student hold the same background beliefs and draw the same inferences. But they still differ in what they are justified in believing. So that difference in justification must be due to different contrast classes being seen as relevant.

Since justification can vary even when evidence does not, it is also possible for a single person to be both justified and not justified at the same time in believing the same claim. For example, the hospital ethics committee might be justified in believing that the doctor's lying would be morally wrong out of a contrast class that does not include moral nihilism, but not justified in believing the same claim out of a contrast class that does include moral nihilism. It would be strange and misleading to say simply that the committee is both justified and not justified. However, it would be just as strange and misleading to say that jumbo shrimp are both large and not large. This conjunction can still be true if it means that jumbo shrimp are large for shrimp but not large for seafood. Analogously, the ethics committee can be both justified and not justified, if this means that it is justified out of a smaller contrast class but not justified out of a larger contrast class.

When a belief is justified out of one contrast class but not another, and both contrast classes seem relevant, then there will be disputes about which contrast class *really* is relevant. A philosopher might say that moral nihilism is really relevant even in ethics committees, so the ethics committee is not really justified. But then a practical person might say that nihilism is not really relevant even in philosophy, so the student is really justified. Is there any way to determine which contrast class is really relevant to a particular belief in a particular context? I doubt it. Admittedly, certain contrast classes seem rele-

vant to some contexts.<sup>47</sup> If people form and join hospital ethics committees in order to reduce immoral practices in hospitals, then it is natural for these people to see moral nihilism as irrelevant. However, a single person can have conflicting purposes and can occupy more than one context at once. Moreover, claims about justification can cross contexts: Imagine a philosophy student who claims that the ethics committee is not justified in believing what it does. Is the student's contrast class (with moral nihilism) or the committee's contrast class (without moral nihilism) the one that is relevant to the student's claim about the committee's belief? And what if the ethics committee claims that the student's belief is justified? When contexts cross in this way, we need to distinguish the person with the belief (the believer) from the person who judges whether the belief is justified (the judger). In judging that a belief is justified or not, the context of the believer seems more important if the purpose of the judger is to predict whether the believer will do well on the committee or in the class. In contrast, the context of the judger seems more important if the judger's purpose is to decide whether to endorse the belief or the process that produced it. Since the judger's purpose affects which context and contrast class seem relevant, there is no way to specify which context or contrast class is *really* relevant for judgments of the form "*S* is justified in believing *p*" even if one adds "when *S* is in context *C*."<sup>48</sup>

Luckily, we don't need to solve this problem here. Regardless of which, if any, contrast class is (really) relevant, we can still talk about justification relative to contrast classes. We can say that the ethics committee is justified relative to a smaller contrast class that does not include moral nihilism, but is not justified relative to a larger contrast class that does include moral nihilism. We can make the same relativized judgments about the philosophy student. And we can understand a speaker who calls a belief justified without qualification because she means that the belief is justified relative to some contrast class that the speaker sees as relevant, regardless of whether that contrast class is really relevant. It can sometimes become important for practical purposes to decide which contrast class is relevant.<sup>49</sup> But judgments about justification relative to a certain contrast class are all we need to describe the epistemological position from the morally neutral standpoint of meta-ethics. Consequently, I will henceforth avoid any talk about what is justified without qualification and any assumptions about which contrast class is relevant. I will talk only about whether a moral belief is justified relative to a specified contrast class.<sup>50</sup>

This approach produces as many levels of justification as there are contrast classes.<sup>51</sup> To be justified relative to {A, K} is different from being justified relative to {A, K, R}, which is different from being justified relative to {A, K, R, N}, and so on. For simplicity, I will focus on two vague contrast classes. As I said, ethics committees usually consider only a limited range of relevant alternatives. The same is true of everyday disputes about important issues, such as abortion or capital punishment. The class that includes all moral positions that would be taken seriously by any ethics committee or any common person will be called *the everyday contrast class*, and I will talk about

*everyday justification* when someone is justified relative to the everyday contrast class. I will not try to specify exactly what falls into this contrast class, because it is very large and indeterminate, and most of its members do not matter here. What is crucial here is that the everyday contrast class does *not* include extreme positions, such as moral nihilism. When two people argue about abortion, they usually assume that either it is morally wrong to get an early abortion or it is morally wrong for someone to physically stop a competent adult from getting an early abortion. They take it for granted that *something* is morally wrong, so they assume that moral nihilism is false.

This assumption is *not* made in some moral philosophy classes, such as mine, that take moral nihilism seriously and ask how it could be refuted. The class of moral positions that are taken seriously in such moral philosophy classes will be called the *philosophy contrast class*, and justification relative to this contrast class will be called *philosophical justification*. Roughly, the philosophy contrast class can be seen as including the everyday contrast class plus moral nihilism (and other extreme views). Because of this relation, everyday justification can be explained in terms of philosophical justification. One is everyday justified in believing *p* if and only if one is philosophically justified in believing the conditional: *p* if moral nihilism is false and if our other everyday assumptions are true.<sup>52</sup>

Because everyday justification assumes so much, it might seem that everyday justification cannot depend on evidence or be tied to truth. But it can. To see this, consider an office betting pool where each worker picks one out of the ten horses in a race. Four horses have already been chosen, and the remaining six include a horse named Playboy. Playboy is better in the mud than the five other remaining horses, and rain is forecast for race day. This evidence can make one justified in believing the conditional that Playboy will win if any of the remaining six horses wins. Nonetheless, if one knows nothing about the four previously chosen horses, or if one knows that one of them is an even better mudder than Playboy, then one is not evidentially justified in believing that Playboy will win out of all ten horses.<sup>53</sup> Analogously, one can have evidence that a certain moral claim is true *if* any moral claim in the everyday contrast class is true, even if one has no evidence for that moral claim out of the whole philosophical contrast class. This happens, for example, if the everyday contrast class includes A, K, and R, the philosophical contrast class also includes N, and one's evidence raises the probability that (if A or K or R, then A) without raising the probability that (if A or K or R or N, then A). One can then be everyday evidentially justified without being philosophically evidentially justified.

Thus, the two distinctions together yield four kinds of justification:

Justification	Everyday	Philosophical
Evidential		
Instrumental		

These distinctions can be generalized and applied outside morality, and more and finer distinctions can be drawn, but these four kinds of justification in moral theory will keep us busy enough.

## Theories of Justification

Now we can ask which kinds of justification can be achieved. I cannot go through every possible theory about justified moral beliefs, but I will briefly discuss intuitionism, coherentism, and contractarianism, because these are three of the most common theories about how moral beliefs can be justified.

### Moral Intuitionism

It seems obvious to me that there is a spot on my pants, because I see one. It also seems obvious that  $2 + 3 = 3 + 2$ . When such beliefs are formed without inferring them from anything else, they can be called noninferential beliefs. Noninferential beliefs are possible, but are they ever justified? Foundationalists in general epistemology assert that some people can be justified in believing some claims even if they do not or could not infer those claims from anything else that they believe.

Moral intuitionists apply foundationalism to moral claims. In a weak version, moral intuitionism asserts that some people are justified in believing some moral claims even though they *do* not infer those claims from any other beliefs. However, people often do not actually go to the trouble to formulate inferences that they could easily make if they just tried. In at least some of these cases, a person's being justified in believing a moral claim still depends on the believer's ability to draw an inference that she does not actually draw, since the person would not be justified in the belief if she could not draw the inference. But even such an ability to infer is not necessary for some moral believers to be justified, according to a stronger version of moral intuitionism. This stronger version will be our topic here, and I will refer to it simply as moral intuitionism.<sup>54</sup>

Moral intuitionism is then the claim that some people are justified in believing some moral claims in some way that does not depend on the believer inferring or even being able to infer the claim from anything else that the person believes. Some moral intuitionists claim that a believer is justified by some nondoxastic state, such as a perceptual state or an emotion, that could not be a premise in an inference.<sup>55</sup> Other moral intuitionists claim that the believer is justified but not by anything at all. In the latter case, the believer is said to be directly justified, and what is believed is often called self-evident. In both cases, the person is said to be justified noninferentially. My arguments will apply to both versions, so I will ask whether believers can be noninferentially justified.

The moral beliefs that are claimed to be noninferentially justified might be particular moral beliefs, such as

(P) It is morally wrong for this doctor to lie to this patient on this occasion about this HIV + test result,

or more general moral principles, such as

(G) It is always morally wrong to lie without an adequate moral reason (or it is always prima facie morally wrong to lie).

Moral intuitionists need not claim that these beliefs are infallible (that is, cannot be false) or are indubitable (that is, cannot be doubted). Nor need they claim that these beliefs are the object of a special faculty of intuition or that one can become justified in believing (P) or (G) just by staring at them. Instead, moral intuitionists can admit that one must know the facts of the case and must reflect on the concepts in (P) or (G) and on applications to other examples in order to be justified in believing (P) or (G).<sup>56</sup> However, if moral intuitionists require some kind of reflection, such reflection must help only to understand the belief and must not provide premises for inferences to the belief. What moral intuitionists deny is that one must be able to infer (P) or (G) from some conceptual analysis, from the facts of the case, from moral beliefs about other cases, from more general moral principles, or from anything in order to be justified in believing (P) or (G).

A moral nihilist would deny both (P) and (G). What can a moral intuitionist say against such a nihilist? Not much. Moral intuitionists can point out that moral nihilism conflicts with many moral beliefs in which many of us feel great confidence after long reflection. However, to appeal to such a moral belief in an argument against moral nihilism begs the question in much the same way as it would beg the question to appeal to a belief about the external world in an argument against Descartes's deceiving demon hypothesis. Some moral beliefs appear obvious, but that appearance is just what would be predicted by the nihilistic hypothesis that all moral beliefs are illusions, so that appearance provides no evidence at all against moral nihilism. Some moral intuitionists might respond that moral nihilists cannot really understand (G), because (G) is a conceptual truth, so anyone who understands it must accept it. However, there is no independent evidence that all moral nihilists must fail to understand (G), so it again begs the question against moral nihilism to assume that (G) is a conceptual truth. Consequently, moral intuitionism cannot provide philosophical evidential justification for any moral belief. Of course, some moral intuitionists would respond that they do not try to provide philosophical evidential justification, and even that it is a mistake to try to provide that kind of justification, but my point is only that they do not and cannot provide philosophical evidential justification.

What about everyday evidential justification? Suppose that a member of the ethics committee admits that some acts are morally wrong, so he or she

is not a moral nihilist, but honestly does not see why lying in this case is morally wrong, as (P) claims, or why lying is always prima facie immoral, as (G) claims. This person can be called a moral deviant, since most people accept (P) and (G). What could a moral intuitionist say in response to this moral deviant? Again, not much. The intuitionist can keep repeating, "Don't you see" or "Reflect some more," but the deviant can honestly try her best and still not agree. The intuitionist might claim that the deviant does not fully understand the disputed claim, but the deviant's understanding is demonstrated by the inferences she draws from her beliefs and by her nondeviant uses of the constituent concepts in other contexts. Then, even if the deviant cannot give any reason against (P) or (G), the deviant still has the same claim to be justified as the intuitionist with the more common beliefs in (P) and (G).

One might think that *both* the deviant and the common moral intuitionist are justified. But consider an analogy. Miss Scarlet and Professor Plum were the only two people in the room when Colonel Mustard was killed. One detective looks at Miss Scarlet and believes that she did it. Another detective looks at Professor Plum and believes that he did it. If neither detective has any evidence, neither seems justified. Both detectives ought to suspend judgment until the opposing hypothesis is ruled out. For the same reason, neither the deviant nor the common moral intuitionist seems justified. Their beliefs might be epistemically permitted, but neither belief has enough positive support to say that anyone ought to believe it instead of the contrary hypothesis.

To get positive support, a moral intuitionist might bring in other beliefs, such as that lying in this case has bad consequences or that other cases of lying are morally wrong. Then some principle of universalizability could be applied. Since a moral deviant accepts some moral judgments, such arguments might give a moral deviant a reason to change her beliefs. However, such arguments also give up moral intuitionism, since a moral belief is not justified *noninferentially* if its justification depends on the ability to infer it from other beliefs.

The only way for moral intuitionists to remain intuitionists is to insist that some moral beliefs are justified without depending on any ability to infer them from any beliefs. But then what *can* make moral beliefs justified? One possibility is that moral beliefs are justified by some mental state that is not and does not imply a belief. Emotions have been said to play this role.<sup>57</sup> But any such mental state could occur in a moral deviant, and the deviant beliefs do not seem justified in the face of the more common contrary beliefs. So such mental states cannot be sufficient by themselves to make a moral belief justified.

A second possibility is that reflection on concepts and examples can make a moral belief justified even if the reflection does not yield anything from which the moral belief can be inferred.<sup>58</sup> However, not just any reflection will do. To make a moral belief justified, reflection must have certain properties: It must be careful and long enough, it must not be distorted by ignorance or self-interest, and so on. Moreover, the moral belief does not seem justified if the believer does not somehow believe that his reflection has these properties.



To see this, imagine someone who says, "I don't think that I reflected long enough, and maybe I didn't reflect on the right kinds of examples, and I might not be impartial or informed, but my reflection still makes me believe that sodomy is immoral." This person does not seem justified in holding this belief. If this belief is not special, then we can generalize to the principle that, if one does *not* believe that one's reflection is adequate or reliable in the circumstances, then one is not justified in trusting the results of the reflection. The problem for intuitionism is then that, if one *does* believe that one's reflection is adequate or reliable in the circumstances, one has all one needs (other than general intelligence) to be able to infer one's moral belief from one's belief about the adequacy or reliability of its source. One might, for example, use an inference like this:

- (i') Moral beliefs based on reflection in circumstances like these are reliable;
- (ii') My moral belief in (P) is based on reflection in these circumstances;
- (iii') Therefore, my moral belief in (P) is reliable (or likely to be true).

A believer need not consciously run through such an inference in order to be justified. Nonetheless, a believer is not justified if she would not see her reflection as adequate or reliable and thus would not accept these premises under appropriate circumstances, so the believer must at least dispositionally believe what is necessary for the inference. Consequently, a moral belief cannot be justified by reflection unless the believer is at least *able* to infer the moral belief from other beliefs. This conclusion is contrary to the central claim of moral intuitionism.<sup>59</sup>

A third possibility is externalism—the view that some beliefs are justified because they have properties that are independent of any other mental state. The most common kind of externalism is reliabilism—the claim that some beliefs are justified because they are reliable (or result from reliable processes). If reliabilism is applied to moral beliefs, then one need not reflect on one's moral beliefs or believe that one's moral beliefs are reliable in order for them to be justified. If one does not believe that one's moral beliefs are reliable, then one cannot *believe* that they are justified, or be *justified* in believing that they are justified, or *show* that they are justified. But one's moral beliefs can still *be* justified as long as they *are* reliable.<sup>60</sup>

This externalist view is consistent, but not convincing. The primary model for reliabilism is vision. However, people usually have evidence that their vision is reliable, since they remember that their visual beliefs in certain circumstances were confirmed in the past by their other senses and by other people. When they saw a cat, they could pet it, and others could see and pet it, too. To test whether reliability *by itself* is sufficient for justification, we need an example in which a believer has no evidence that the belief source is reliable under the circumstances.<sup>61</sup> Imagine that Leslie notices an aura around Joan's head. She hasn't seen anything like this before, and she can't

understand why she sees it now, but there seems to be a faint brown tinge where some strange kind of energy seems to be emanating from Joan's head. Leslie considers three hypotheses: either aural energy is emanating, or it is a visual illusion, or there is something wrong with Leslie's eyes. Leslie believes that aural energy is emanating, but is she justified in believing this (instead of the other hypotheses)? It seems not. What if there really are auras, and she is a reliable detector of them? She still does not yet seem justified in believing in auras. One reason is that Leslie need not even believe that she is reliable or justified. If we ask her whether she is any good at detecting auras, she might say, "I have no idea. Maybe; maybe not." That response would make it hard to see her belief as justified. Even if Leslie does believe in her own reliability, she has no evidence for her beliefs or for her reliability or against the other hypotheses. If she continues to see auras and meets other people who see similar auras, and if auras help to explain independent phenomena, then Leslie might become justified in believing in auras. But prior to such evidence, she is not justified in believing in auras.

The reason is *not* that Leslie might be wrong and is not certain. Mistakes are possible even after she gains evidence for her reliability, but her fallibility and uncertainty at that later stage do not keep her from being justified then. The point is also *not* that Leslie has evidence against auras. I assume that auras are compatible with Leslie's beliefs and past experience, and even with science (at least because scientists should initially be open to new forms of energy). And I am *not* saying that Leslie is irrational and should stop believing in auras. Her belief might be permitted epistemically. Nonetheless, Leslie still does not have enough positive support to be justified in believing in auras.

Leslie's aura beliefs are much like noninferentially justified moral beliefs according to reliabilism. If one does not believe and has no reason to believe that one's moral beliefs (or moral belief forming processes) are reliable, then one's moral beliefs are no more justified than Leslie's beliefs about auras, even if one's moral beliefs turn out to be reliable. Of course, most people have experienced agreement with other people's moral beliefs, as well as connections among their moral beliefs and their other beliefs, so their moral beliefs still might be justified on that basis. But in the absence of such other beliefs, moral beliefs seem no more justified than Leslie's aura beliefs, even if both are reliable.

Of course, much more could be said for and against reliabilism, but not here. I hope that I have said enough to suggest why reliability is not sufficient by itself to make moral beliefs justified. I argued earlier that emotions and reflection are also not enough to make moral beliefs justified. There might be other kinds of noninferential justification, but I suspect that they will run into similar problems. So moral intuitionism is wrong about everyday evidential justification for moral beliefs like (P) and (G).

This conclusion assumes that a moral belief like (P) or (G) can be seriously questioned in an everyday context, such as a hospital ethics committee. But what about the belief that it is morally wrong to torture babies just for fun? A view that denied this belief would never be taken seriously in any everyday

context. The everyday contrast class was defined to include only those views that would be taken seriously in an everyday context. So, if anything in the everyday contrast class is true, then it is morally wrong to torture babies just for fun. This moral belief is thereby everyday evidentially justified, even if it is not inferred from anything.

However, this is just an artifact of what everyday justification is. To say that a belief is justified at the everyday level is in effect to say that it is justified given our everyday assumptions. It should not be surprising that our everyday assumptions are justified given themselves. That triviality does not even begin to show that such beliefs are true or likely to be true. All it shows is that these beliefs are assumed to be true. Moreover, even this kind of everyday justification will apply only to moral beliefs that are extremely weak. If any normal person would have any serious doubts about a noninferential moral belief, then that belief will lack even this kind of everyday evidential justification.<sup>62</sup> The only ways to avoid all everyday doubts are either to qualify the moral belief so that its application will be in dispute (such as "It is morally wrong to harm for an inadequate reason") or to limit it so much that it does not apply to many, if any, real acts (such as "It is morally wrong to torture babies just for fun"). Consequently, no serious dispute will be resolved by claiming that a moral belief is justified noninferentially.

Moral intuitionism still might provide *instrumental* justification for a wider range of moral beliefs. Whether a belief is justified instrumentally depends on the effects of being in that belief state. If I strongly believe (P) or (G) after careful reflection, then to give up these beliefs might be costly to my self-confidence, my psychological health, and my personal relations, regardless of whether I can justify these beliefs inferentially. If so, the very factors that moral intuitionists cite as justifying some moral beliefs *can* instrumentally justify some moral beliefs. This instrumental justification holds even when moral beliefs cannot be inferred from other beliefs, and even when moral beliefs are contrasted with moral nihilism.<sup>63</sup> So some noninferential moral beliefs are instrumentally justified at both the everyday level and the philosophical level.

There can also be instrumental justification for believing moral intuitionism, that is, for believing that some moral beliefs are justified noninferentially. If one does not believe that one's belief in (P) or (G) is justified noninferentially, then one might think that one needs to infer it from something. But justifying inferences are often hard to find, and their premises are even harder to justify. In fact, one could not form any moral belief if one first had to justify every assumption inferentially. And it does not seem to be worth the trouble to test a moral belief by means of inference in trivial cases where everyone agrees already. So there can be instrumental justification not only for a moral belief but also for the belief that a moral belief is justified noninferentially.

This might explain why many people believe in moral intuitionism, but the point should not be overestimated. Instrumental justification still has no tendency to show that the justified beliefs are true or probably true. Thus,

the extent of the justification that is captured by moral intuitionism is very limited. It can be summarized in following diagram:

Justification	Everyday	Philosophical
Evidential	Only trivial cases	None
Instrumental	Some	Some

### Moral Coherentism

The next common theory of justification to be considered is coherentism. Coherentists claim that a person *S* is justified in believing *p* to the extent that *p* coheres with a coherent system of beliefs held by *S*. A system of beliefs is coherent to the extent that its beliefs are jointly consistent, comprehensive, and connected by logical, probabilistic, and explanatory relations.<sup>64</sup> And a particular belief coheres with a system of beliefs to the extent that the belief has these relations to the system. When *p* does bear these relations to *S*'s other beliefs, *S* has all it takes (other than general intelligence) to be able to infer *p* from another belief. Thus, coherentists claim that *S* is never justified in believing *p* unless *S* can justify *p* by some inference. Coherentists need not require that *S* explicitly formulate any inference, but what makes *S* justified also makes an inference available. Coherentists also need not require that *S* be able to connect *p* to every other belief in *S*'s system, or that *S* have a completely comprehensive or consistent system of beliefs. This ideal is impossible in practice. Nonetheless, one system of beliefs can still be more coherent than another, and the more coherent *S*'s system of beliefs is, the more justified *S* is in believing what coheres with that system. Coherentists can also hold that some minimal degree of coherence is necessary for a system to provide enough positive support to make *S* justified at all.

The most straightforward application of coherentism to morality occurs when people try to connect disparate moral beliefs about particular cases by inferring them from a general moral principle. This goal is called narrow reflective equilibrium.<sup>65</sup> For a simple example, suppose *S* believes:

- (a) It is morally wrong for a student to tell a professor that his dog ate his paper, when this is not true, just to avoid a penalty for late papers;
- (b) It is morally wrong for a professor to tell a student that chapter 6 will be on the exam, when it won't, just to make this student study more;
- (c) It is morally wrong for a professor to tell a student that chapter 6 will not be on the exam, when it will, just to keep this student from becoming too depressed when he tries to study this difficult chapter.

These beliefs and others might seem to suggest that lying is always wrong, but

- (d) It is not morally wrong to lie to prevent a murder.

Beliefs (a)–(d) can all be inferred from a general principle that

- (G) It is always morally wrong to lie without an adequate reason,

together with subsidiary beliefs, such as

- (c') Preventing depression is not an adequate reason to lie;  
(d') Preventing murder is an adequate reason to lie.

One can then also infer a particular conclusion in the lying doctor example:

- (P) It is morally wrong for this doctor to lie to this patient on this occasion about this HIV + test result,

because this act comes under (G) and (c'), assuming the doctor's only reason to lie is to prevent the patient's depression. Of course, many more beliefs and principles would have to be added to make this system comprehensive, but then *S* is supposed to be justified in believing (P) because of how (P) coheres with the rest of the system, and *S* is also justified in believing each of the other claims because of how it coheres with the remaining beliefs. This structure is circular, because the justification of (a) refers to (b), and the justification of (b) refers to (a), but this circle is not supposed to ruin the justification if the system is comprehensive enough.

However, this theory runs into several problems. First, it is not clear whether anybody actually has a system of moral beliefs that is comprehensive or connected or even consistent, or that will remain consistent when enough moral beliefs are added to make it comprehensive. If nobody has a coherent enough system of moral beliefs, then no moral beliefs are justified by the standards of coherentism.

Second, even if someone's moral system is coherent, there will always be incompatible systems that are equally coherent. In fact, if one wants to reject any particular conclusion, one can always deny that conclusion and modify one's principles to form another system that is just as consistent, connected, and comprehensive as the original system. For example, if one wants to avoid the conclusion (P), one can keep (a)–(d), (G), and (d') but replace (c') with something like

- (c'') Preventing depression is not an adequate reason to lie, except when the depression would ruin the rest of someone's life before he dies of a deadly disease.

Of course, some people might see (c'') as ad hoc or implausible, but the system with (c'') but not (P) is just as consistent, connected, and comprehensive as the system with (c') and (P). More particular moral beliefs might be added to argue against (c''), but a defender of (c'') could either deny these moral claims or incorporate them by redefining terms or by adding qualifications or other principles. In the end, there will always be alternative systems of moral beliefs that are equally coherent but conflict in their beliefs about (P).<sup>66</sup>

To solve this problem of underdetermination, Rawls and his followers<sup>67</sup> add that justified moral beliefs must cohere not only with substantive moral principles but also with social science and with ideals of the person and of society. This wide reflective equilibrium might somewhat reduce the indeterminacy of narrow coherence. Nonetheless, uncontroversial findings of social science rarely, if ever, suffice to rule out a competing moral view; and ideals of the person and society are usually just as hard to justify as the beliefs that they are supposed to justify. For these reasons, even wide reflective equilibrium cannot avoid competing, equally coherent systems.

When their coherence is equal, neither system can be justified in contrast with the others on the basis of coherence alone. *S* might already hold the beliefs in one of the systems, or the beliefs in one system might seem more credible to *S*, but that cannot make *S* justified in believing that system in contrast with the other, unless *S*'s prior beliefs or feelings of credibility are somehow justified independently.<sup>68</sup> This justification cannot come from coherence (since the systems are equally coherent) or from moral intuition (because of the arguments in the section on moral intuitionism), so there is as yet no way for *S* to be justified in believing one of the equally coherent systems in contrast with the other.

In addition to underdetermination, coherentism also faces a problem of truth. Coherence provides evidential justification only if a coherent system is more likely to be true, but coherence requires only relations among beliefs inside a system, so how could internal coherence increase the probability of truth?

In response to this problem in general epistemology, many coherentists introduce a new order of evidence.<sup>69</sup> The most common kind of evidence for a first-order belief (such as, that my house is on fire) is another first-order belief (such as, that smoke is pouring out of it). A different kind of evidence cites second-order beliefs about first-order beliefs. In this model, one uses some first-order beliefs to justify second-order beliefs about the circumstances under which certain first-order beliefs are reliable, and then these second-order beliefs are used as evidence for other first-order beliefs. For example, past experience teaches us that we are not reliable at identifying persons by sight when the light is bad, the person is far away, we do not know the person well, and so on; but in the absence of these conditions, visual identifications seem reliable. These beliefs plus my present visual experience then imply that it is probably my wife whom I see next to me. Because such second-order beliefs imply that certain first-order beliefs are reliable or likely to be true,

if all of these beliefs form a coherent system, then its coherence might be a reason to believe that the system is likely to be true.

This approach can be extended to moral beliefs.<sup>70</sup> Evidence for beliefs about what is morally right or wrong usually takes the form of other beliefs about what is morally right or wrong. However, one can also cite second-order beliefs about such first-order moral beliefs. For example, one might find that many of one's first-order moral beliefs that were formed when one's self-interest was involved or when one did not fully understand a situation had to be rejected later because they conflicted with many other beliefs. This kind of consideration might give members of an ethics committee reason to discount as unreliable the deviant moral beliefs of a doctor who wants to avoid the unpleasant task of telling a patient that he or she is HIV +. It might also give them a reason to discount the moral beliefs of people who do not understand important aspects of the medical situation.

If one had a complete list of distorting factors, then one might infer that one's moral beliefs are reliable when all of those factors are missing. But many problems arise. First, it would be hard to tell whether one's list of distorting factors is complete. If even one factor is overlooked, it might be the one that distorts one's moral belief. Second, it is not clear whether some factors belong on the list. Is someone's moral view of abortion more reliable or less reliable if they are passionate about the issue? Third, even if we have a complete list of distorting factors, it is often hard to detect these factors in particular moral beliefs. It is usually easy to tell whether my vision is operating in good light, at short distances, and so on, but it is hard to tell whether a moral believer is adequately informed or impartial. Did the hospital ethics committee miss some relevant fact? Is their self-interest involved in some roundabout way (such as that they will want to lie to one of their patients in the future, or will themselves be patients in the future)? Moral beliefs are almost never formed with complete information; they almost always affect one's welfare somehow; and we can't tell when ignorance or partiality is operating; so there is always some reason to doubt their reliability.<sup>71</sup> Finally, even if a moral belief lacks all of the factors that make moral beliefs *unreliable*, this does not prove that the moral belief is reliable. Our moral beliefs might not work well even under the best conditions.

Nonetheless, if one finds that, in the absence of a certain range of factors (such as self-interest, ignorance, and irrationality), one's moral beliefs rarely need to be revised, are widely accepted by other people, and support and explain one another, then all of this together is *some* reason to believe in the reliability of moral beliefs formed under such circumstances. One cannot ever be sure that one did not overlook some crucial fact or that one's self-interest is not operating below the surface, but one can still have some reason to believe that a moral belief is not distorted in such ways. In such a situation, one can have some reason to believe that a particular moral belief is true. The belief still *might* be false but that does not remove the reason to believe that it is true, since justification can be fallible.

Does this kind of coherence provide everyday evidential justification? I think that it *can* in some cases. In some cases, all of the alternatives that would be seriously considered in everyday contexts conflict with moral beliefs formed under circumstances that usually lead to moral beliefs that are reliable according to other commonly shared moral beliefs, and there is no support for those alternatives except beliefs that were formed in circumstances that usually distort moral beliefs, again according to commonly shared moral beliefs. This might happen, for example, when all of the members of a hospital ethics committee agree about a case where they do not know any of the parties involved, and do not themselves have to face similar situations, but do take the time to study the case carefully. Even those who disagree, such as the doctor who does not want to tell a patient that he or she is HIV +, might have enough other moral beliefs to require the doctor to admit that his or her own belief in this case is probably distorted. When all of this comes together in the right way, coherence methods can provide everyday evidential justification. This method can work even in some cases where moral beliefs were controversial before being tested by coherence, so it can achieve more than appeals to intuition.

Of course, there will be many other cases where everyday conflicting moral beliefs cannot be dismissed as distorted, and where one does not have adequate reason to believe in the reliability of one's own moral beliefs. Then one's moral beliefs are not everyday evidentially justified. I suspect that most of our moral beliefs fall into this class, because of the problems mentioned above. But the fact that the coherence method fails in most or even almost all cases does not show that coherence cannot yield everyday evidential justification in some other cases.

It is important, however, not to read too much into this claim. Everyday justification just assumes in a way that moral nihilism is false (along with other extreme views); so to have everyday evidential justification that lying is wrong is just to have philosophical evidential justification for believing the conditional that lying is wrong *if* anything is (and if our everyday assumptions are true). Everyday justification can thus be seen as a kind of conditional philosophical justification.

What about *unconditional philosophical* evidential justification? To achieve this kind of justification, coherence arguments would have to work against moral nihilism. However, the coherence method needs to use some first-order moral beliefs to justify second-order beliefs about conditions of reliability. Moral nihilists would deny those first-order moral beliefs, along with claims about their reliability. Thus, this coherence method begs the question against moral nihilism.<sup>72</sup> The problem is that the coherence method depends on shared moral beliefs, but moral nihilists do not share any moral beliefs, so the coherence method cannot even get started. This is analogous to the problems in using coherence among empirical beliefs to refute the hypothesis of a deceiving demon. In both areas, coherence can provide at most everyday but not philosophical evidential justification.<sup>73</sup>

What about *instrumental* justification? The main question here is whether coherence with a coherent system of beliefs can make a moral belief state have good effects. It seems so. When a moral belief is connected in certain ways to one's other beliefs, one cannot give up the moral belief and maintain a coherent system without changing many other beliefs. It is sometimes a good thing to shake up one's beliefs, but large changes in one's belief system often have detrimental effects on one's confidence, plans, relations to other people, and so on. When too much change would cause such problems without compensating benefits, one can be instrumentally justified in holding on to a moral belief because of its coherence with one's other beliefs. The same points apply to extreme views. Most people could not accept moral nihilism without radical and disruptive changes in their belief systems, so coherence can make a moral belief instrumentally justified at the philosophical level as well as the everyday level.

There can also be instrumental justification for testing moral beliefs by their coherence with one's other beliefs. Coherence testing takes time, but it can reduce the risk of having to change one's mind later, which can be disruptive in some settings, such as hospital ethics committees. Moreover, people often object when an ethics committee's decisions in different areas appear arbitrary or incompatible, but such objections can be reduced if the committee's various decisions are seen as connected under general principles. For such reasons, the ethics committee and the hospital will tend to run more smoothly if they try to form their moral beliefs into a coherent system. This method can be used improperly, but sometimes it has better effects than not checking for coherence. When this is so, one can be justified instrumentally in checking one's beliefs for coherence.<sup>74</sup>

Nonetheless, this instrumental justification still does not show that coherent moral beliefs are true or even likely to be true. There are reasons to check for coherence and not to give up beliefs that do cohere, but the kinds of justification that are provided by coherence are still limited. They can be diagrammed as follows:

Justification	Everyday	Philosophical
Evidential	Some nontrivial cases	None
Instrumental	Some	Some

### Contractarianism

The third common method for justifying moral beliefs is contractarianism. Actually, many different methods go by the name of contractarianism. Often a contractarian framework is used only as a way to display connections among prior moral beliefs, and then coherence is the basic justification.<sup>75</sup> In contrast, other contractarians claim that a moral belief is justified if certain people under certain circumstances would agree to the moral belief, or to a rule or system that yields it, independent of whether the result coheres with prior

moral beliefs.<sup>76</sup> It is this latter version of contractarianism that concerns us here, because this version claims that agreement is an independent source of justification.

In contrast with moral intuitionism and coherentism, contractarianism is primarily about when moral beliefs are justified, not about when people are justified in believing them. Most people do not become justified in believing moral claims by arguing or even being able to argue that certain people would agree on any beliefs or rules. Yet if contractarians are right, a person can be justified in holding a moral belief if he or she is justified in believing that the moral belief is justified by some agreement. Moreover, contractarianism need not apply to all areas of morality. Most contractarian theories are mainly about justice and not about benevolence or other virtues. Nonetheless, if any such agreement can make anyone justified in believing any moral claim, this would be enough to refute moral skepticism.

But who must agree in order for a moral belief to be justified? One possibility is that a moral belief or rule is justified because all rational people would accept it. If this kind of argument worked, it would accomplish a lot. Rationality is normative, but it is supposedly not moral, so this argument would justify moral beliefs by reference to nonmoral norms. And if all rational people would accept a certain moral belief or rule, then moral nihilists must be irrational, since they do not accept the moral belief or rule. So this kind of contractarian argument would provide not just everyday but also philosophical justification if it worked.

Of course, this kind of argument will provide justification only to the extent that one is justified in accepting its theory of rationality, instead of one of the many competing theories. Moreover, on any plausible nonmoral theory of rationality,<sup>77</sup> it is hard to see why all rational people must agree on any definite moral belief or rule. The most common examples are moral rules against killing and breaking promises. The basic idea is that, if people in society did not generally accept and follow these rules, then society would break down, which would harm everyone, so every rational person would accept these rules for society. However, this agreement disappears when the rules are defined more precisely. Consider a more limited rule against killing people who are over six feet tall. Shorter people would not accept this rule as the only rule against killing, and they are rational, so not all rational people would accept this limited rule as the only rule against killing. And shorter people might refuse to cooperate with anyone who does not accept a rule that protects them as well. Nonetheless, it still might not be irrational for some people over six feet tall to reject the unlimited rule "Don't kill anyone" in favor of the more limited rule "Don't kill anyone over six feet tall," if these taller people don't have much to fear or to gain from shorter people. If so, neither rule or set of rules must be accepted by all rational people, so neither is justified in this way.

Moral contractarians admit this, so they add other constraints. They claim that a moral belief or rule is justified when it would be accepted by all people who are both rational and impartial. Or they claim that all rational people must

agree under certain circumstances (usually tailored to ensure impartiality) or that all rational people need to cooperate with everyone in society. Contractarians often also add constraints on the form of the rules to be chosen.<sup>78</sup>

However, these new constraints create new problems. First, why should a moral skeptic or nihilist accept these constraints? The demand for impartiality and the claim that certain circumstances are appropriate for a social contract themselves need to be justified. It is hard to see how such constraints could be justified without depending on more basic moral beliefs about the nature or content of morality. But that would beg the question against moral nihilism and other extreme views. So such contractarian arguments do not seem to provide philosophical justification.

Even for everyday justification, there is a problem of multiple contracts, similar to the problem of multiple coherent systems. Philosophers have developed many different accounts of rationality, of impartiality, and of the circumstances that are appropriate for the agreement that is supposed to justify moral rules or beliefs. Different accounts of rationality, impartiality, and appropriate circumstances imply that rational impartial persons under appropriate circumstances would agree on different moral beliefs or rules. Thus, one can use the contractarian framework to justify many different moral systems. Simply adjust one's accounts of rationality, impartiality, and appropriate circumstances to fit one's favorite moral system. If the method is flexible enough to accommodate two competing moral systems, then it cannot justify one moral system in contrast with the other. This problem can be solved only by justifying particular theories of rationality, impartiality, and appropriate circumstances that pick out a single moral system. But it is hard to see how this could be done (especially for impartiality and appropriate circumstances) without depending on moral beliefs that beg the question against competing moral systems in the everyday contrast class. So such contractarian arguments do not seem to provide everyday justification for any controversial moral belief.<sup>79</sup>

Even if these problems could be solved, a more basic problem would still arise. There are two kinds of rationality. Sometimes a person is called rational to the extent that she adjusts her beliefs to the evidence. If everyone who is rational in this way would accept a moral belief, then that moral belief would be evidentially justified. However, one could not justify a claim that all people who are rational in this way would accept a certain moral belief if one did not already have separate evidence for that belief. The claim about agreement would not add any new evidence.

In any case, contractarians do not use this evidential kind of rationality. They see people as rational to the extent that they seek good effects and avoid bad effects of some kind. But then, if all rational people accept a belief or rule, this must be because accepting it has good effects or avoids bad effects. And such effects do not show that a moral belief is true or even likely to be true. To see this, suppose that someone will kill anyone who does not believe in Santa Claus, or that anyone without this belief will suffer debilitating depression. Then everyone who is rational in the relevant way would accept the belief in Santa Claus. But none of this provides evidence for Santa Claus.

Similarly, to show that all rational people would accept certain moral beliefs as ways to protect or improve their lives or society does not show that these beliefs are likely to be true or that they are justified evidentially. Instead, contractarian arguments provide only instrumental justification.

Contractarians usually respond that their arguments do not rest on contingent preferences that some people lack and also cannot be used to argue for the opposite moral beliefs. Even if so, this still does not make rational agreement evidence of moral truth. Suppose that a belief in Santa Claus is necessary to prevent everyone from becoming so depressed that they would lose all motivation to do anything, regardless of what they prefer to do. And suppose that no opposite belief could provide any motivation. That still would not be evidence for Santa Claus. In general, the effects of a mental state cannot show that the content of the mental state is likely to be true. So, even if contractarian arguments do not depend on specific preferences and do not work for contrary beliefs, they still provide at most instrumental and not evidential justification.

Another response is that contractarian arguments are about effects of moral rules or standards rather than of beliefs. Rules and standards have the form of imperatives (for example, "Don't kill"), so they have no truth value, and one should not expect them to be evidentially justified. One can still be evidentially justified in believing that killing is morally wrong if this moral claim just means that killing violates a rule or standard that is justified by the effects of accepting it.<sup>80</sup> But is this really what the moral claim means? I doubt it. The claim that an act is morally wrong seems to be about the act and not about the effects of accepting a rule against the act. People can form such moral beliefs without thinking at all about rules or their effects. Moreover, a rule can be beneficial in certain circumstances even though the corresponding moral belief is not true. Suppose that an evil demon will kill everyone who does not accept the rule "Do not dance." Then this rule would be accepted by all rational people, and even by all rational impartial people, so that rule would be justified, according to contractarians. But this is not evidence that it is morally wrong to dance.<sup>81</sup> This example is fanciful, but similar problems arise in common cases. For example, some opponents of abortion admit that society would be better off if people did not accept "Don't kill fetuses" (or at least that not all rational people must accept "Don't kill fetuses"), but they still believe that abortion is morally wrong, because it is unfair to kill fetuses to benefit others. Whether or not they are right, their position is at least consistent. And the same goes for anyone who believes that a harmless act, such as sodomy, is immoral. This shows that "morally wrong" does not mean the same as "violates a rule that is justified by its effects" (or by agreement among everyone who is rational in a way defined by effects). Consequently, moral beliefs are not evidentially justified just because the corresponding moral rules are instrumentally justified.

Contractarians might respond that such noncontractarian beliefs are not really about morality, because the only real morality is defined by rational (impartial) agreement based on the effects of accepting rules. However, that appeal to definition is no better than saying that non-Euclidean geometry is

not really about geometry because it does not share Euclidean axioms. Such definitions just change the subject. If all that contractarians mean when they say that abortion or sodomy is not morally wrong is that not all rational (impartial) people would accept rules against these acts, then opponents of abortion can admit that those acts are not wrong on contractarian definitions, or that they are not contractarian-wrong. This does not imply that these acts are not morally wrong in the sense claimed by opponents of abortion and sodomy, and denied by others. Because the original issue was not about rational (impartial) agreement, contractarians cannot resolve that issue by defining moral wrongness in terms of their own theory. Contractarians still might provide evidence for something else, but they do not provide evidence for morality.<sup>82</sup>

Contractarians could, of course, give many more responses, and there are many more kinds of contractarianism. Nonetheless, the basic point remains and can be generalized. Because contractarian arguments show at most that certain moral rules or the corresponding moral beliefs would or should be accepted because of their effects, contractarians do not show that the contents of those moral beliefs are likely to be true. So contractarian arguments do not provide evidential justification. Nonetheless, contractarians still might succeed in showing that certain moral beliefs have better effects in certain circumstances than any everyday alternative, and even than moral nihilism. This would show that some moral beliefs are instrumentally justified at both the philosophical and everyday levels. If this much is granted, then the kinds of justification that contractarians provide can be diagrammed as follows:

Justification	Everyday	Philosophical
Evidential	None	None
Instrumental	Some	Some

### Conclusion

All of these theories deserve much more attention than I have been able to give them here. There are also many other theories about justified moral beliefs (including pragmatism, contextualism, derivations of *ought* from *is* or from moral language, and so on) and many other attempts to rule out moral nihilism. I cannot discuss all such theories and attempts here.<sup>83</sup> Nonetheless, I hope that my discussion of the three most common theories has suggested some reasons why these other theories also cannot provide philosophical evidential justification, even if they can provide some everyday justification and some instrumental justification. If so, we can never have philosophical evidential justification, but we can still get some everyday evidential justification for some moral beliefs, as well as instrumental justification at both the everyday and philosophical levels. Thus, some kinds of justification can be obtained, but others cannot.

This conclusion illuminates the arguments for moral skepticism. Skeptical hypothesis arguments assume that we are not justified in our moral beliefs if we cannot rule out skeptical hypotheses, such as moral nihilism. I have argued that this claim is true if it is about philosophical justification. However, we can still be everyday justified in ignoring a skeptical hypothesis, because, if one chooses out of the everyday contrast class, one must choose an alternative that implies the denial of the skeptical hypothesis, since they all do.<sup>84</sup> Thus, the skeptical hypothesis argument shows that our moral beliefs cannot be philosophically evidentially justified, but it does not show that our moral beliefs are never everyday justified.

A similar conclusion can be reached regarding the regress argument for moral skepticism. Premise (2) denied noninferential justification, and premise (7) denied circular justification. But I argued that moral intuitionism and coherentism can supply everyday evidential justification for some moral beliefs. If so, premises (2) and (7) are false if they are about everyday justification. But they are still true of philosophical evidential justification, because moral intuitionism and coherentism do not provide philosophical evidential justification, as I also argued.

I conclude that moral skepticism is right about philosophical evidential justification and wrong about everyday evidential justification and about instrumental justification. I call this conclusion *limited moral skepticism*, because it denies a kind of justification that is important to those of us who want to understand the limits on our human condition. Although everyday evidential justification can be obtained, all this means, as I said, is that we can be philosophically justified in believing that certain acts are morally wrong *if anything is*, since everyday justification simply ignores moral nihilism. Still, some might think that even some everyday evidential justification is enough, so they might see my conclusion as a denial of moral skepticism. I do not care. Whether my view is classified as moral skepticism depends on which contrast class one takes to be relevant and to define moral skepticism. I already said that I will not try to decide which contrast class is most important or relevant. That varies with context and purpose. In the general discussion here, it is enough to determine exactly which contrast classes one can be justified relative to, and how.

### Notes

For many very helpful comments, I thank Robert Audi, Malia Brink, Bob Fogelin, Bernie Gert, Mitch Haney, Richard Hare, Stephen Jacobson, Lawrence Kim, John Konkle, Chris Kulp, Paul McNamara, Jim Moor, Peter Railton, Stefan Sencerz, John Skorupski, David Sosa, Bill Throop, Mark Timmons, Jon Tresan, Doug Weber, Nick Zangwill, and audiences at University College Dublin, Trinity College Dublin, Queen's University of Belfast, the Universities of Edinburgh, Glasgow, and St. Andrews, and the Moral Epistemology Conference at Dartmouth College.

1. See Edmund Gettier, "Is Justified True Belief Knowledge?" *Analysis* 23 (1963), 121-3.

2. See William Alston, "Justification and Knowledge," reprinted in *Epistemic Justification* (Ithaca: Cornell University Press, 1989), 172–82.
3. To avoid convoluted phrasing, I will sometimes write about whether a moral belief is justified, but this should be taken as shorthand for whether a believer is justified in believing a moral claim, except where otherwise indicated. In any case, my main points will apply to both notions.
4. "Noncognitivism" is misleading, since etymology suggests that cognitivism is about cognition or knowledge. There are implications for knowledge (see below), but linguistic moral skepticism is directly about truth values.
5. See my "Some Problems for Gibbard's Norm-Expressivism," *Philosophical Studies* 69 (1993), 297–313; and Peter Railton in the volume.
6. See Simon Blackburn in this volume. Some emotivists and prescriptivists who accept only minimal truth in moral claims still might deny moral knowledge if they take knowledge to imply some stronger kind of truth. Then their views still imply moral knowledge skepticism.
7. See J. L. Mackie, *Ethics* (Harmondsworth: Penguin, 1977), 30–42. Error theorists do not see *all* moral claims as false; since, for example, "Abortion is not morally wrong" is true if "Abortion is morally wrong" is false. That is why I restrict the thesis of error theories to positive moral claims.
8. Philosophers also debate whether there is always a distinctively moral (as opposed to self-interested) reason to be moral, and whether there is always enough reason to make it irrational to be immoral or only to make it not irrational to be moral. These debates yield distinct kinds of practical moral skepticism.
9. Throughout, qualifications like "just to . . ." and "just for . . ." are meant to refer not to the agent's motives but only to the absence of any other reason for the act.
10. Here I disagree with David Copp, "Moral Skepticism," *Philosophical Studies* 62 (1991), 203–33.
11. This last argument derives from Gilbert Harman, *The Nature of Morality* (New York: Oxford University Press, 1977), 3–10, although Harman is not a moral skeptic.
12. *Outlines of Pyrrhonism*, book I, chapter 15.
13. See Robert Audi in this volume and my section on moral intuitionism below. The more general claim that some beliefs are noninferentially justified is called "foundationalism." General foundationalists need not embrace moral intuitionism, since they might deny that any moral beliefs are among the noninferentially justified foundations.
14. Thelma and Louise cannot both be noninferentially justified if noninferential justification entails truth or if it entails that anyone who understands the claim believes it. So the skeptical argument would have to be reformulated to apply to such accounts of noninferential justification. The point would then be that neither Thelma nor Louise would have any reason to believe that she is the one who is noninferentially justified.
15. One well-known example is John Searle, "How to Derive 'Ought' from 'Is,'" *Philosophical Review* 73 (1964), 43–58, reprinted with responses and related articles in W. D. Hudson, ed., *The Is/Ought Question* (New York: St. Martin's, 1969).
16. For these premises to be descriptive, the relevant F, C, G, and R must be defined without reference to any norms. The relevant group, for example, cannot be the group of mature or normal or rational people.
17. See Alan Gewirth, "The 'Is-Ought' Problem Resolved?" reprinted in *Human Rights* (Chicago: University of Chicago Press, 1982); and *Reason and Morality* (Chicago: University of Chicago Press, 1978).
18. Many sociobiologists are skeptics, but some draw substantive moral conclusions from sociobiology. See Richmond Campbell, "Sociobiology and the Possibility of

- Ethical Naturalism" in *Morality, Reason, and Truth*, eds. David Copp and David Zimmerman (Totowa, N.J.: Rowman and Allanheld, 1984), pp. 270–96.
19. See Lawrence Kohlberg, *The Philosophy of Moral Development* (San Francisco: Harper and Row, 1981), chapter 4.
20. See Richard Brandt in this volume, and his *A Theory of the Good and the Right* (Oxford: Clarendon Press, 1979).
21. See Christopher Morris and David Copp in this volume; David Gauthier, *Morals by Agreement* (New York: Oxford University Press, 1986); and Bernard Gert, *Morality* (New York: Oxford University Press, 1988). R. M. Hare in this volume might also be classified in this group, since his justification rests on nonmoral normative claims about rationality and about how we should use moral language.
22. More limited forms of contextualism, such as that developed by Mark Timmons in this volume, might not be subject to the objections in the text, but then it is not clear that they are incompatible with moral skepticism.
23. See Geoffrey Sayre-McCord in this volume; David Brink, *Moral Realism and the Foundations of Ethics* (Cambridge: Cambridge University Press, 1989), chapter 5; and my section on moral coherentism below.
24. See Brink, *Moral Realism and the Foundations of Ethics*, appendix 1.
25. First Meditation, in *Meditations on First Philosophy*.
26. To rule out a hypothesis is to give a good reason not to believe that hypothesis. For example, one can rule out a hypothesis by showing that it is internally inconsistent or that it is incompatible with the facts, although the facts must be established independently so as to avoid begging the question.
- In order to be justified, it is *not* necessary to go through any actual *process* of ruling out any hypothesis or to *show* that any hypotheses is false. All that is necessary is that one have the *ability* to rule out certain hypotheses, and that requires only the justified beliefs and intellectual abilities needed to rule out those hypotheses. Consequently, my argument does not conflate the state of being justified with the process of justifying or with showing that one is justified.
27. *Pace* arguments that are sometimes attributed to Hilary Putnam, *Reason, Truth, and History* (New York: Cambridge University Press, 1981), chapter 1.
28. Moral nihilism implies that nothing is morally required, because an act is morally required only if not doing it is morally wrong. A broader moral nihilism would also deny that anything is morally good or bad, morally should or ought to be done, and similarly for all other predicates. This more general hypothesis would be needed to argue for skepticism about these other kinds of moral claims, but the arguments would be basically the same, so I will focus on moral wrongness.
- Even with this focus, moral nihilism is distinct from (a) the modal claim that nothing *could* be morally wrong and (b) the second-order claim that no first-order claim of the form "x is morally wrong" is *true*. Moral nihilism on my definition is (c) the negative existential claim that there does not exist anything that is morally wrong. This quantifier does not range only over acts that have actually been done. Instead, (c) denies any substantive claim that any act or act type is morally wrong. Although (c) is then implied by (a) and (b), (c) still does not imply (a) or (b), because (c) is not about what is possible or about truth values.
29. I just assume these claims here, but I will argue for them in "Skepticism and Nihilism about Moral Obligations" *Utilitas* 7 (1995), 217–36.
30. Some might want to add that it must also be true that *p* entails *q*, or that I am justified in believing the conjunction of *p* and (*p* entails *q*), but these additions would not affect the main points here.



31. Parallel arguments can be constructed with other extreme hypotheses, such as Moral Egoism (nothing is morally wrong unless it is against the agent's interest or, alternatively, nothing is morally wrong if it is in the agent's interest), and Moral Crackism (nothing is morally wrong except stepping on cracks in sidewalks).

32. This kind of response is associated with Wittgenstein and in ethics with Philippa Foot, *Virtues and Vices* (Berkeley: University of California Press, 1978). Compare Simon Blackburn in this volume. Moral constructivism is also sometimes said to make moral nihilism incoherent and thus to rebut moral skepticism, much as phenomenalism rebuts skepticism about the external world. See Brink, *Moral Realism and the Foundations of Ethics*, pp. 19–20 and 31–5. However, if any such theory makes moral nihilism incoherent, that theory will need to be defended.

33. Hannah is also not justified in believing that the box contains a game instead of a puzzle just because either (a) she does not consciously think of the possibility of a puzzle, or (b) she has no positive reason to believe that the box contains a puzzle, or (c) her acts would not be affected if the puzzle hypothesis were true. Thus, not being considered, not being supported by a reason, and not affecting one's acts do not make a hypothesis irrelevant; and this applies to skeptical hypotheses as well.

34. If it is epistemically irresponsible to believe anything without positive support, then a belief might be permissively justified only if it is positively justified. But this antecedent is at least controversial, so permissive and positive justification should be distinguished initially as concepts.

35. How much is enough? That seems to depend on context, which might also affect what counts as support at all. I do not need to resolve these controversial issues here, since my arguments do not depend on requiring any unusual degree of support. In particular, my kind of moral skepticism does not require justified moral beliefs to be certain, as is often claimed by critics of skepticism.

36. Some moral epistemologists ask whether certain factors, such as intuition or coherence, ever provide any positive support that would be adequate if it were neither undermined nor defeated. A moral belief supported to this extent is then said to be *prima facie* justified. However, this weak thesis is compatible with the skeptical view that no moral belief is ever justified adequately, since support for moral beliefs is always undermined or defeated. To address such moral skepticism, I will focus on the question of whether factors, such as intuition or coherence, ever make a belief justified by providing adequate positive support.

37. My distinction derives from David Copp, "Moral Skepticism." Compare also Gregory Kavka, "The Toxin Puzzle," *Analysis* 43 (1983), 33–36.

38. Some philosophers suggest that evidence and evidential justification need only "aim" at truth, but skeptics need not deny that people aim at truth in ways that do not increase their chances of obtaining it, so they need not deny that people can be evidentially justified in the weak sense that requires only aiming at truth.

39. This definition does not require the support to be a belief, so it does not beg the question against moral intuitionists who see some beliefs as justified independently of any other belief.

40. Jon Tresan pointed out that one might call a belief justified because a belief state has value in itself apart from its effects. This kind of justification would not be instrumental, but it also would not provide evidence.

41. There are exceptions, such as that it is more likely that one will be happy if one believes that one will be happy. See William James, "The Will to Believe," *New World* (1896), 327–47. But such special cases do not include moral beliefs, because moral beliefs are not about their own effects.

42. A critic might argue that instrumental justification is not epistemic if it is not tied to truth. Even if so, instrumental justification is still important to moral epistemology, because justifications for moral beliefs are often instrumental. See the section on contractarianism in this chapter.

43. Linguistic moral skeptics and error theorists must deny that moral beliefs can be justified evidentially, because they deny that moral beliefs can be true, but they can still see some moral beliefs as justified instrumentally.

44. On contrast classes in epistemology, see Fred Dretske, "Epistemic Operators," *Journal of Philosophy* 67 (1970), 1007–1023; and Alvin Goldman, "Discrimination and Perceptual Knowledge," *Journal of Philosophy* 73 (1976), 771–91. On contrast classes in explanations, see Lipton and van Fraassen in *Explanation*, ed. D. Ruben (Oxford: Oxford University Press, 1993). I am deeply indebted to these works. The justification of actions is also relative to contrast classes, since I can be justified in paying my debt as opposed to not paying it, even if I am not justified in paying it today instead of yesterday (when it was due) or in paying it with a five dollar bill instead of with five one dollar bills (for which I have no reason). I suspect that all reasons (whether explanatory or justificatory, for acts or beliefs) are relative to contrast classes, but that broader thesis is not essential to my position here.

45. Some scholars deny that Kant was an absolutist, but here I am discussing the absolutist view that is often ascribed (rightly or not) to Kant.

46. Contrast classes do not resolve all disagreements. Ethics committee members can still disagree about whether a belief is justified relative to the same contrast class.

47. Pragmatists and contextualists, such as Mark Timmons in this volume, try to work out the factors that affect which contrast class is (taken to be) relevant; and feminist critiques of authority, such as by Margaret Walker in this volume, ask who gets to say what will be (seen as) relevant. These projects thus complement mine.

48. This suggests that there is no way to settle disputes about whether a believer is justified without qualification, unless judges happen to share the same purpose. This might lead to another kind of moral skepticism.

49. For example, a government body might have to decide who to follow when protestant hospital ethics committees recommend something that Catholic hospital ethics committees dismiss out of hand. But moral epistemologists need not take sides. We can say that Catholics are justified relative to a Catholic contrast class even if they are not justified relative to a larger Protestant contrast class. This is all we need to say if our goal is to understand the dispute rather than to settle it.

50. Critics might object that these relativized judgments are not normative. However, normative judgments are often relativized, such as when I call certain apples "good for this time of year." See also note 44 on paying bills. Moreover, even if relativized judgments are not essentially normative, they can still be used normatively (that is, for normative speech acts) on particular occasions by people who assume that certain contrast classes are important.

51. I assume that a contrast class must include the belief that is justified (or else one could not believe it *out of* the class) and at least one other belief (or else there would be no contrast) and that all members of the class conflict in some way (to avoid irrelevancies). Actually, the members of a contrast class will be sets of beliefs, but for simplicity I will often speak as if they are single beliefs.

52. I am *not* saying that *S* is everyday justified in believing *p* if and only if it is true that: if nihilism and other extreme positions are false, then *S* is philosophically justified in believing *p*. This cannot be right, because then one could detach "*S* is philosophically justified in believing *p*" when nihilism and other extreme positions are

false and *S* is everyday justified. My claim is instead that *S* is everyday justified in believing *p* if and only if *S* is philosophically justified in believing the conditional: if nihilism and other extreme positions are false, then *p*. This allows one to detach "S is philosophically justified in believing *p*" only when one has the premise "S is philosophically justified in believing that nihilism and other extreme positions are false." I am indebted to John Skorupski for forcing me to clarify this.

My claim is also *not* about justification relative to a contrast class that includes general skeptical hypotheses, such as Descartes's deceiving demon. One might not be justified in believing *any* conditional relative to that unlimited class. But the philosophical contrast class at issue here includes only moral positions that are taken seriously in moral philosophy classes, where general skeptical hypotheses are not considered and many nonmoral beliefs are assumed. The relevant conditional can be justified relative to this moral philosophy contrast class because it is just a claim about evidential relations independent of any substantive moral views. I am indebted to Mark Timmons for pressing me on this point.

53. In a more theoretical analogy, if there are three competing theories of astrology, and two are inconsistent with certain observations, but the third is consistent with those observations, then those observations are evidence that, if any of these theories is true, the third theory is; so one can be evidentially justified in believing the third theory out of this limited contrast class. Judgments like these enable us to make sense out of what astrologers say and do.

54. The main reason to focus on the strong version of moral intuitionism is that skeptical arguments apply if even an ability to infer is required. Notice also that moral coherentists need not deny the weak version of moral intuitionism if they do not require justified believers to formulate the inferences that constitute coherence. There might still be disagreements about when an inference is available and when one is able to infer a belief, but I will not try to pin down these slippery notions here.

55. This claim is sometimes called a moral sense theory and distinguished from moral intuitionism, but I use "moral intuitionism" broadly to include moral sense theories and even moral reliabilism (discussed below).

56. See Robert Audi in this volume.

57. See William Tolhurst, "On the Epistemic Value of Moral Experience," *Southern Journal of Philosophy*, Supplement to Volume 29 (1990), 67–87, and my response, "Moral Experience and Justification," in the same issue.

58. See Robert Audi in this volume.

59. This argument is derived from Lawrence Bonjour, *The Structure of Empirical Knowledge* (Cambridge: Harvard University Press, 1985), section 2.3; and Brink, *Moral Realism and the Foundations of Ethics*, chapter 5, section 5.

60. Some externalists admit that one is not justified in believing (P) if one also believes that this belief is *not* reliable. But all they require is the absence of belief in unreliability. If the person who believes (P) does not have any beliefs about whether or not this belief is reliable, this belief is still supposed to be justified if it is reliable.

61. Such examples are given in Bonjour, *The Structure of Empirical Knowledge*, section 3.2, and Keith Lehrer, *Theory of Knowledge* (Boulder, Colo.: Westview Press, 1990), 163.

62. Although the everyday contrast class was defined to include whatever might be seriously considered in any everyday context, the range of views that would be taken seriously in a particular context is often much more limited. More moral beliefs can be noninferentially justified relative to these more limited contrast classes, but this is still only because those beliefs are assumed and not disputed in such contexts.

63. Jon Tresan pointed out that an instrumental justification would beg the question against moral nihilism if it assumed a value that moral nihilism denies, but instrumental justifications can avoid this problem by using nonmoral values.

64. This account of coherence derives from Geoffrey Sayre-McCord, "Coherence and Models for Moral Theorizing," *Pacific Philosophical Quarterly* 66 (1985), 170–90. See also his essay in this volume.

65. By Norman Daniels, "Wide Reflective Equilibrium and Theory Acceptance in Ethics," *Journal of Philosophy* 76 (1979), 256–82.

66. This is an extension of Nelson Goodman's new riddle of induction in *Fact, Fiction, and Forecast* (Indianapolis: Hackett, 1979), 72–83. In terms of contrast classes, Goodman's problem can be partly resolved by saying that past experience justifies us in believing that emeralds are green as opposed to blue but does not justify us in believing that emeralds are green as opposed to grue.

67. See John Rawls, *A Theory of Justice* (Cambridge: Harvard University Press, 1971); and Daniels, "Wide Reflective Equilibrium and Theory Acceptance in Ethics."

68. The problem of credibility is presented forcefully by Stefan Sencerz, "Moral Intuitions and Justification in Ethics," *Philosophical Studies* 50 (1986), 77–95.

69. Especially Bonjour, in *The Structure of Empirical Knowledge*, who builds on the work of Wilfrid Sellars.

70. See Brink, *Moral Realism and the Foundations of Ethics*, chapter 5, sections 7–8.

71. More reasons to doubt reliability in moral beliefs are given by Margaret Walker in this volume when she shows how many moral beliefs depend on dubious relations of power and authority, and by some recent work in cognitive psychology on common mistakes in moral beliefs. See Daniel Kahneman's Tanner lecture, "The Cognitive Psychology of Consequences and Moral Intuition" (forthcoming).

72. The illegitimacy of using nonshared moral beliefs in coherence arguments is admitted by Rawls, "Outline of a Decision Procedure for Ethics," *Philosophical Review* 60 (1951), 183; Brink, *Moral Realism and the Foundations of Ethics*, p. 132; and Daniels, "Wide Reflective Equilibrium and Theory Acceptance in Ethics," p. 259, and "Reflective Equilibrium and Archimedean Points," *Canadian Journal of Philosophy* 10 (1980), 85–89. These coherentists do not, however, see the relevance of this "independence constraint" to moral nihilism.

73. Coherence might provide philosophical evidential justification if one could construct a meta-justification as in Bonjour, *The Structure of Empirical Knowledge*, section 8.3. However, even if this kind of meta-justification works for empirical beliefs, which is controversial, nothing along these lines seems to work for moral beliefs.

74. In addition, if one checks a moral belief against as many other beliefs as possible, one has done as much as one can and as much as can be expected. One cannot then be faulted, and one is permissively justified. Nonetheless, permissive justification does not show that one's moral belief is true or likely to be true. Even if coherence is the best available tool for forming moral beliefs, it still might have no relation to truth, because even the best tools might fail most of the time.

75. See John Rawls, *A Theory of Justice*, p. 21.

76. See David Gauthier, *Morals by Agreement*, p. 269.

77. One could just define "rational" so that a certain moral system would be chosen. For example, one could say that people are rational only if they maximize utility for all, and then all rational people would agree on utilitarianism of some kind. But that just begs the question against alternative moral systems and beliefs.

78. For example, Rawls requires a "veil of ignorance" to ensure impartiality, and

he adds formal constraints, in *A Theory of Justice*, sections 23–24. For criticisms and alternatives, see chapters 5 and 13 of Gert, *Morality*.

79. Some moral beliefs (such as that it is morally wrong to torture babies just for fun) might be justified in any contractarian framework that would be acceptable to any person in any everyday context. Such beliefs would be everyday justified, but this is just because they reflect common everyday assumptions, so this kind of justification would not show that these beliefs are true.

80. See David Copp in this volume.

81. A contractarian might respond that dancing is immoral if it leads the demon to kill someone. However, the demon in my example would not kill anyone for dancing as long as they still accept the rule against dancing, so that they feel bad about dancing. In this situation, accepting the rule is instrumentally justified, but the corresponding moral belief is still not evidentially justified.

82. Contractarian theories might seem to be about morality because their conclusions overlap to some extent with common moral beliefs. I doubt that this is the right way to tell which views are about morality. After all, religious rules often overlap in content with moral rules, but religion is still separate from morality. And moral nihilism is about morality, even though its content does not overlap with common moral beliefs. A better way to determine which views are about morality is to ask which claims are seen as consistent. Moral nihilism is not seen as consistent with moral wrongness, but denials of contractarian-wrongness are seen as consistent with morality. This suggests that moral nihilism is a moral view, but denials of contractarian-wrongness are not moral views.

83. I will discuss some of the other theories in my *Moral Epistemology* (Englewood Cliffs, N.J.: Prentice-Hall, forthcoming).

84. It might also seem that premise (3'), the principle of closure, fails. However, it appears to fail only when its antecedent refers to justification relative to a smaller contrast class than in its consequent. Closure seems to hold as long as the contrast class remains constant, and that is all that should be expected of the principle.