

7

Ought to Is

The Puzzle of Moral Science

John Basl and Christian Coons

7.1 INTRODUCTION

Our moral commitments influence our views about the empirical world. Consider, for instance, the strong correlations between political ideology and beliefs about climate change, the president's religion/birthplace, or the effects of gun control.¹ We presume these correlations are partly explained by our propensity to see the world in ways that comport with our moral outlook. This seems a poor way of arriving at beliefs. It threatens the common ground for dialogue between those with different moral outlooks; but more directly, our evaluations (or their contents) seemingly cannot count as evidence for claims about the earth's climate, the president's birthplace, or the effects of gun control. As obvious as this may seem, the puzzle is how to best explain why this is so. Thus, here is our puzzle:

The Puzzle: What justifies, or at least explains, our resistance to adopting our views about the empirical world in light of our evaluative or moral views.

In short, we want to know what's wrong with inferring from ought to is, or at least why we balk at such inferences. While many doubt that purely descriptive premises can deductively entail any moral conclusion, deduction from ought to is seems relatively straightforward. For example, if ought entails can, then (i) *we ought to reduce net suffering* entails (ii) *we can reduce net suffering*. Sometimes our evaluative views logically entail empirical claims—and yet when they do it nevertheless seems illegitimate to draw

¹ See Kahan and Braman 2005; McCright and Dunlap 2011; Jones 2010; Crawford and Bhatia 2012.

the conclusion. For example, consider the following inference, call it *Ought to Is*:

Ought to Is

- A. It is never wrong to do what maximizes utility.
- B. It is always morally wrong to kill an innocent child.
- C. Therefore, killing an innocent child never maximizes utility.

A and B are, of course, moral claims, C is a descriptive generalization, and A and B appear to logically entail C. Interestingly, however, it also appears to be a mistake to *infer* C from A and B.

Yet diagnosing the mistake is difficult. To help illustrate, imagine a clever student in an introductory ethics class—call him “Smart Alec”—who says the following:

I’ve reflected on last week’s arguments for utilitarianism, and I found them pretty darn convincing. While I’m still not sure about the view, I am confident that it’s at least always permissible to maximize utility—after all, what could be wrong about doing what would be impartially best for everyone? And I’ve always felt confident—just from reflecting on the act itself—that killing an innocent child is always wrong. So, thanks to this class, I’ve had to rethink many of my assumptions, and I’ve learned a lot. Specifically, I now see that killing an innocent child actually never will maximize utility. Moral reasoning is much more powerful than I’d initially thought!

Alec’s comment sounds like a joke; but what’s so funny? One is inclined to instruct Alec to avoid the inference and revisit A or B. But on what grounds? The premises *are* consistent. And surely no one can claim to have observed cases where such killing did, in fact, maximize total utility. And it won’t be very powerful to reply—“But remember the hypothetical example involving the time-bomb? You know, the one with the rocket sled and the stadium full of children who are even more innocent. We all agreed killing the innocent child in this case *would* maximize utility—I stipulated that it would—and you agreed!” To this, Alec could simply reply—“Yes, but I now see that situations like that don’t actually occur.” And for all we know, maybe they don’t.

Even though the *type* of evidence Alex has for A presumably differs from B, this isn’t problematic—one may infer from premises whose mode of justification differs. And though Alec’s belief in A and B may not be fully justified, unless one is willing simply to reject the predominant modes of gathering moral evidence, or adopt the view that justification *never* transfers over known entailment, Alec appears to have at least some justification for both A and B. So why doesn’t he have at least *some* (even if defeasible and weaker) justification for C? What are we to say to Alec? More generally, given that moral claims can have empirical entailments, why doesn’t moral reasoning serve as a tool for learning about the physical world? On its face,

this kind of research, call it “moral science,” seems not only less reliable, but wholly illegitimate. Our chapter assumes this is true and attempts to explain why.

We will begin by outlining and then rejecting the two extant approaches to this puzzle; both are advanced in Alex Barber’s (2013) “Science’s Immunity to Moral Refutation.”² Each approach, Barber warns, comes at a price. The first—call it the “anti-realist strategy”—requires the bold assumption that we are all, often unwittingly or implicitly, moral anti-realists. The price of his second approach—we call it the “realist-friendly strategy”—is that it relies on an intolerably restrictive moral epistemology. Specifically, it assumes that no *particular* moral claims are justified unless they are deduced from *a priori* moral principles and empirical background assumptions. We argue that we needn’t pay either price because neither strategy plausibly explains the puzzle.

Next, we’ll briefly canvas a number of tempting alternative approaches to resolving the puzzle—we’ll argue each is a dead end. Learning what we can from these failed solutions, we advance our own proposal. We argue that every attempted moral science inference faces a dilemma. The moral scientist takes all the premises of such inferences to be true “come what may” (i.e. true regardless of how the empirical facts turn out to be) or they do not. If they do not, we will argue that the moral scientist must employ a false premise or rely on non-moral assumptions that already entail the conclusion. On the other horn, if the moral scientist takes all the premises of their inference to be true come what may, their premises entail that a known contingent proposition is necessary. Given this knowledge, they ought to reject at least one of their premises.

7.2 BARBER’S ANTI-REALIST STRATEGIES

Barber’s (2013) “Science’s Immunity to Moral Refutation” aims to do more than explain why we resist inferences from ought to is; he also claims our resistance anchors a powerful case for moral anti-realism—a case he calls the “argument from moral immunity” (633). Specifically, Barber claims that

² Barber’s focus, however, is slightly different than our own—focusing on why moral premises can’t overturn or count against a scientifically supported conclusion. But we stress, and Barber seems to acknowledge, that the phenomenon is broader: not only do we tend to think “moral science” can’t be used to revise our scientifically supported beliefs, it also shouldn’t even count as a reasonable way to form beliefs (e.g. inferring C in *Ought to Is*) when we don’t already have scientific evidence for (or against) the conclusion. The difference is of little import as most of the explanations Barber considers, if successful, also answer our puzzle.

anti-realists of various stripes can “explain easily” why inferring from ought to is is illegitimate, while realists, he says, can do so only by placing “significant constraints on the structure of evidence for moral judgments” (633). Anti-realism’s unique ability neatly to resolve our puzzle, he concludes, is a powerful reason to think that our moral discourse is implicitly anti-realist.

We will examine why Barber thinks realists struggle with our puzzle in Section 7.3; here we focus on his case that “each of the four main types of anti-realism”—non-cognitivism, subjectivism, skepticism, and error theory—can easily explain why it’s illegitimate to infer from ought to is. Barber’s arguments in this section are relatively short, so we’ll present them using his own text:

Moral non-cognitivism: Non-cognitivists (“expressivists”) hold that our moral utterances express, not genuine judgments of belief, but mental states for which truth is an inappropriate measure of evaluation. Premises that are not even truth-apt, let alone true, cannot legitimately be used to infer anything that is truth-apt; hence we can never reasonably infer from our moral convictions to the falsity of a scientific theory with which they seem to conflict. (637)

Moral subjectivism: Moral subjectivists allow that moral utterances express genuine judgments of belief but deny that these judgments are objective. Their truth conditions are tied too closely to the inclination of the judger to make the judgment. If this is right, the feature of our moral convictions that explains why we cannot draw on them to overturn scientific claims is poverty of content. Their content is so proximal that they cannot be used to draw objective scientific—distal—conclusions. (638)

Moral skepticism: Skepticism deserves to be classified as a form of moral anti-realism only as it tends towards extreme and generalized pessimism about the quality of evidence available in the moral sphere. Sceptics of this calibre can readily account for its being wrong to pitch moral considerations against scientific evidence, since an argument is only as secure as its premises. (638)

Error theory: Error theorists have the simplest explanation of all for why we should not call on our cherished moral beliefs to settle a scientific question: they aren’t true. (638)

Barber knows none of these views, even if true, could explain the puzzle by themselves. The puzzle, after all, is how to vindicate *our* widespread rejection of inferences from “ought” to “is.” So while these anti-realist views may entail the illegitimacy of moral science, they could only explain the puzzle insofar as *we*—the population at large—actually accept these views and sense their implications for moral science. Remarkably, that’s just what Barber suggests. Specifically, he proposes that we reject these inferences because “below the radar of conscious reflection we are all moral anti-realists” (638).

The anti-realist strategy seems rather incredible. Most of us are not avowed anti-realists, so it postulates a collective delusion about our own

meta-moral views. Barber provides no independent evidence for such a delusion, but he carefully acknowledges and addresses this worry, and denies that it's damning. Perhaps it isn't. We do, after all, have many implicit beliefs, and we may be implicitly sensitive to the underlying rules or structure of our moral discourse while being wildly inaccurate when describing it explicitly. And so, if elements of our practice are really best explained by implicit anti-realism, perhaps we are indeed implicit anti-realists.

The primary problem with Barber's anti-realist strategy, however, isn't its worrisome appeal to implicit belief; instead, it is that the implicit beliefs would each explain "too much"—if we really had these anti-realist attitudes we would reject other forms of reasoning that we find cogent. Consequently, though the anti-realist strategy does, in one sense, "easily" explain our puzzle—in that it is sufficient for the explanandum—we'll contend it's not a plausible explanation. Plausible explanations, of course, also cohere with our other background assumptions, especially those related to the phenomena in question. So though the hypothesis that *Americans fear wheeled vehicles* is similarly sufficient for the explanandum *Americans don't bike to work*, it does not cohere well with our love of cars and mild fondness for recreational biking. Each of the anti-realist strategies, we'll argue, are similarly in tension with our broader background views—specifically our other views about what we can infer from, and what is entailed by, moral claims.

First, consider the proposals that widespread and implicit acceptance of skepticism or error theory explains the puzzle. Subconscious acceptance of skepticism or error theory shouldn't only assert itself in our hesitancy to *infer* from is to ought, it would also predict our hesitancy to accept any substantive moral claim—whether it's in conflict with our empirical judgments or not. But we don't balk at moral premises; we balk at certain inferences *from* them. Moreover, such a picture conflicts with our tendency to claim that some substantive moral claims are *more* justified than others. So, if implicit anti-realism explains our rejection of moral science, the implicit theory does not seem to be error theory or skepticism.

Two things may be worth noting before we move on to the next anti-realist strategy. First, this is in no way an argument against moral error theory or skepticism. Rather, we are merely pointing out that our discourse does not behave as one would expect were we implicitly to accept either view. Second, for all we've said, it remains a remote possibility that we actually are implicitly anti-realist, and our being so somehow explains the puzzle. But for that possibility to become defensible, plausible auxiliary hypotheses are needed to explain why, for example, we don't object to moral premises generally and often debate their relative merits. Until such hypotheses are advanced, we think it's safe—dialectically at least—to deny that implicit skepticism or error theory explains the puzzle.

Barber's subjectivist strategy is intriguing. Though it's not characterized very precisely, the rough idea appears to be that for subjectivists the content of moral claims are facts about our own attitudes; these "proximal" contents can't license inferences to any general and/or attitude-independent scientific conclusions—i.e. particular facts like "I resent what you did" seem incapable of playing any useful role in any inferences to any general empirical conclusion.

This strategy again proves too much, but unlike error theory and skepticism it also proves *too little*. First, it proves too little because it predicts that we'll find it less problematic to infer from ought to is when the "is" is highly proximal and lacks the generality of a *scientific* conclusion. This prediction fails: our hesitancy to infer from ought to is also extends to *particular* empirical conclusions, even about ourselves, and not just general scientific hypotheses. For example, suppose I reason about what I ought to do now and reach a particular moral verdict: *I must save Sally*. While it's contentious, this verdict arguably entails that *I physically and psychologically can save Sally*. These certainly appear to be quite "proximal" in Barber's intended sense, but even here we don't treat our judgments of duty as grounds for inferring that we can physically and psychologically do our duty. My verdict that I ought to save Sally isn't a reason to believe that I can reach her in time, or that I'm not in a Frankfurt case (etc.); instead we tend to think of these as presuppositions of my verdict. In sum, the subjectivist strategy appears to falsely predict that our concern about inferring from ought to is will wane provided the conclusion is about our own psychology, but it does not.

Barber might respond that even a highly particularized non-moral conclusion such as *I'm psychologically capable of intending to save Sally now* is still somehow "too distal" to be inferred from the relatively proximal content of a (subjectivist) moral claim. Indeed, it's hard to assess our objection because it's not yet clear what, precisely, the subjectivist answer to the puzzle is supposed to be. But there's a second, more serious, problem with the subjectivist strategy no matter how it is fleshed out. Unlike the skeptical or error-theoretic strategies discussed above, the subjectivist strategy attacks the *validity* of inferences from ought to is. But challenges to the validity of the inferences will, like the skeptical or error-theoretic strategies, prove *too much*. The problem is that inferring from ought to is is actually legitimate—perhaps even rationally *required*—when the premises are certain or hypothetically stipulated. For while we want to criticize Alec for believing that (C) *killing an innocent child will never maximize utility* in virtue of his beliefs that (A) *it is never wrong to do what maximizes utility* and (B) *it is always morally wrong to kill an innocent child*, denying that C is entailed by A and B is also a mistake. Were we to ask someone "Would C have to be true if A and B were true?" we'd be equally baffled by someone who answered

“no”—yet this is precisely what the subjectivist strategy suggests one should say. Indeed, when Alec infers C from A and B, we’re not tempted to see his reasoning as a non-sequitur. On the contrary, it’s *because* A and B entail C that we insist that he revisit either A or B. Roughly, we’re inclined to point out to Alec that A and B are inconsistent unless C is true. And because he has no reason to accept C, he must reconsider A or B. Of course, the puzzle itself is explaining why he has no reason to accept C, but the best answer to the puzzle should preserve our thought that someone like Alec is under pressure to revise his premises, so denying the validity of the inference won’t do.

These concerns also beset Barber’s final anti-realist strategy: implicit non-cognitivism. According to this strategy, we implicitly hold that moral utterances express mental states that are not truth evaluable, and as such, it’s natural for us to also think that moral science’s “premises” can’t support any truth-apt conclusions (2013, 637).³

This explanation “explains too much” in two respects. First, if we implicitly and generally accept that moral claims are not truth-apt, then we should also expect at least puzzlement about truth-preserving inferences using moral and non-moral premises to *moral* conclusions, and puzzlement about how arguments using “embedded” or unasserted moral claims could be valid.⁴ For example, it’s puzzling how an argument like (i) *if stealing causes unhappiness, then stealing is wrong*, (ii) *stealing causes unhappiness*, thus (iii) *stealing is wrong* could seem obviously valid, if we’re implicit non-cognitivists. To be clear, we acknowledge that philosophers might one day explain (or may have already explained) how a non-cognitive theory could adequately capture the validity of these inferences; but the fact remains that explaining the validity of these inferences is (or was) an immediate puzzle for non-cognitivists—something that they acknowledged, *prima facie*, looks like a problematic implication of their views. But ordinary speakers are not puzzled; for them these arguments look just as valid as any analogue without moral premises. And so, if ordinary speakers accept non-cognitivism, we would expect them to be puzzled or initially balk at these inferences too—but famously, we do not.

In any case, implicit non-cognitivism fails plausibly to explain the puzzle for a separate reason already mentioned in this section: some inferences from *ought* to *is* seem legitimate—perhaps rationally *required*—when the premises

³ Barber is aware that this explanation is simplistic and that only some versions of non-cognitivism can explain the puzzle. For example, it’s not clear whether Quasi-Realist accounts and/or those that also are non-cognitivists about epistemic norms have the resources to address the puzzle in the manner Barber suggests (2013, 637).

⁴ A *locus classicus* here is Geach (1965).

are known, certain, or stipulated. When one is entitled to certainty about the premises in *Ought to Is*, for example, then one *may* (and seemingly rationally *must*) infer that it never maximizes utility to kill an innocent child. Implicit non-cognitivism, as characterized by Barber, can't explain why the inference suddenly becomes legitimate when we are permitted to assume the premises *are true*. Indeed, if we implicitly accepted non-cognitivism we'd find further mystery (or apparent category error) in the very notions of *knowing*, *being certain of*, or *stipulating*, some moral claim. In short, an inference from moral premises, like in *Ought to Is*, can be legitimate in some theoretical contexts and Barber's conjecture that we're implicit non-cognitivists won't account for this—for it appeals to the content of moral claims (or lack thereof), a feature that won't vary across the relevant contexts.⁵

Thus, the general problem with Barber's anti-realist solutions is that they explain the phenomena but are—barring further development—poor explanations. Nevertheless, the approach may be worth revisiting if there are no plausible alternatives. Accordingly, we turn to Barber's realist-friendly strategy, and subsequently we will advance our own.

7.3 BARBER'S REALIST-FRIENDLY APPROACH

After canvassing and powerfully rejecting grounds a realist might give for rejecting moral science, Barber suggests the realist may have only one defensible option, one that requires a “strategic division” of moral claims into two types. However, Barber emphasizes that the solution comes at a steep price: the strategic division (as the name suggests) may be *ad hoc*, and requires realists to “deny that particular moral judgments about either fictional scenarios or real events have any independent evidential support” (2013, 652). Next we will outline what the strategic division is, how it's supposed to explain the puzzle, and then show that it fails; it cannot, for example, explain what's wrong with our framing inference *Ought to Is*.

According to Barber, the realist must approach the puzzle by dividing our moral judgments into two classes—“M-class” and “m-class”—distinguished by their epistemic character. M-class judgments are subject only to *a priori* justification, immune to *a posteriori* support, while members of the m-class, in contrast, are derived from judgments in the M-class in conjunction with

⁵ For an extended case for why the non-cognitivist must treat these inferences as irrational (and hence couldn't explain their rationality in privileged epistemic positions), see Dorr (2002). Barber acknowledges this exception to the illegitimacy of ought to is inferences later in his paper (2013, 645), but never addresses its deep tension with the “implicit non-cognitivism” approach that he advances.

some *a posteriori* background assumptions. These m-class judgments, by stipulation, cannot themselves enjoy any independent support beyond their derivability from the M-class in conjunction with these background assumptions.

Given this division, Barber argues that a realist could maintain that all inferences from ought to is employ at least one m-class judgment. Thereby, he claims, the justification of these premises will depend on the same considerations that serve as evidence for or against the conclusion. In his words, “whenever we have a moral argument against a (real or putative) scientific finding, one of the moral premises will depend for all its plausibility on the argument’s conclusion, making the argument essentially circular” (Barber 2013, 645).

To illustrate, consider the following inference:

Permissible to Painless

- i. My eating fish is morally permissible.
- ii. If eating fish caused them severe and avoidable pain, then eating them would not be permissible.
- iii. Therefore, my consumption of fish does not cause fish severe and avoidable pain.

This inference does not seem merely fishy, a diagnosis is readily at hand. The first premise is surely an m-class judgment; typically someone who accepts (i) will acknowledge that the permissibility of eating fish depends not only on the badness of pain but also, crucially, on *a posteriori* background assumptions about both whether fish feel pain and the necessity of consuming fish.

As we discuss in Section 7.5, there is a problem of this sort for some but not all moral science inferences. For now let’s assume Barber is right: *any* inference from ought to is that includes an m-class judgment is circular in the manner exemplified in *Permissible to Painless*. Even so, Barber’s strategy can resolve the puzzle only if all the problematic inferences from ought to is employ such judgments. Call the hypothesis that all such inferences employ an m-class judgment the Little-m (L-m) Hypothesis.

For a significant class of moral judgments, it is impossible to tell in advance whether these judgments belong to the M- or m-class. This is because whether a judgment belongs to one class or the other will often depend on which normative theory is true. For example, if Rule Utilitarianism is true “lying is wrong” is presumably *a posteriori*, and hence an m-class judgment, whereas on some Deontological theories, like Kant’s, this may be an M-class judgment. Absent full moral knowledge, one couldn’t tell whether judgments that might serve as the premises of inferences from ought to is were of the M- or m-class and so, at best, there’s no way to assess the L-m Hypothesis.

Even worse for Barber's realist solution, the L-m Hypothesis is, by our lights, simply false. Consider again the inference used to motivate our puzzle:

Ought to Is

- A. It is never wrong to do what maximizes utility.
- B. It is always wrong to kill an innocent child.
- C. Therefore, it never maximizes utility to kill an innocent child.

The conclusion of *Ought to Is* is empirical and yet recall that Alec's reasoning needn't be remotely empirical or *a posteriori*. He may reach the first premise because it seems inconceivable to him that *doing the most good can be wrong* and inconceivable to him that *something could be good, but not good for*. And, like many before him, he might support the second by claiming that grasping the nature of *killing an innocent* itself justifies concluding that it's always wrong.⁶ It seems, at least from the perspective of Alec, that this inference involves only M-class judgments; if this is so, the L-m Hypothesis must be false.

It isn't clear what resources Barber has to respond to the problems facing the L-m Hypothesis. Are we to accept, without some argument, that Alec can't reason to these premises in these ways? That one of his judgments must belong to the m-class? To do so would, somewhat ironically, simply beg the question in favor of the L-m Hypothesis and thereby the generality of the realist-friendly solution. Barber's strategy rests on the undefended assumption that no empirical conclusions could follow from *a priori* moral principles. This assumption looks like it's false. What we need instead is an explanation of why *a priori* principles like Alec's can't be jointly justified, and ideally it would also be nice to be able to explain what's wrong with inferences like *Permissible to Painless* without accepting the strategic division's highly restrictive—"it's all deduction from principles"—moral epistemology. We'll attempt to do just that in Section 7.5, but before doing so, it will be instructive first to consider some tempting solutions that we think are ultimately dead ends.

7.4 ALTERNATIVES THAT ARE NOT

Initially, several alternative approaches to the puzzle seem tempting. Barber's paper includes an excellent survey of attractive but unsuccessful strategies

⁶ Of course, it's tempting to think that both of these premises can't be *a priori* justified. Many of us are inclined to think that having *a priori* justification for one of these premises serves as a reason to discount the other. Of course, whether we may do so depends on a resolution to the puzzle.

(2013, 639–44); here he effectively disarms the proposals that the puzzle can be explained as (1) a corollary of the autonomy of ethics, (2) an effect of our relatively weak evidence for moral claims, or (3) an implication of moral supervenience. To avoid redundancy with his work we won't rehearse those arguments here. Instead, we will strengthen Barber's case that there's no easy answer to the puzzle by considering and rejecting three further superficially plausible alternatives.

7.4.1 Metaphysical Priority Entails Epistemic Priority?

It may be tempting to begin with the hunch that our epistemic puzzle is explained by a *metaphysical* relationship between “is” and “ought.” Though the details are disputed, there's a rough consensus that particular moral facts depend in some sense or other on the non-moral ones—acts have the moral status they do *because* or in *virtue of* their intrinsic and/or relational non-moral features. Perhaps this metaphysical priority, whatever it amounts to precisely, may explain why the non-moral cannot be inferred from the moral.⁷

This tempting idea, by itself, offers no solution; we often reasonably infer independent from the dependent—and this appears to be true across types of dependence relations, e.g. causal, explanatory, or constitutive. For example, wholes depend on their parts, but we regularly infer from wholes to parts—consider our inferences from objects we directly observe to conclusions about the unobservable parts that compose them. Similarly, the dispositions of an object in a fixed context will depend on the object's intrinsic features, and yet the following inference type seems perfectly reasonable:

Dispositional to Intrinsic

1. This object has disposition d in context c.
2. Only things with intrinsic properties X, Y, or Z have d in c.
3. Therefore, this object has X, Y, or Z.

Here, like inferences from wholes to parts, inferring 3 from 1 and 2 is generally faultless. So, typically at least, there is no problem with moving from premises regarding the dependent to the independent; indeed, abductive reasoning characteristically takes just this form.

One might insist there's a special problem with inferences from dependent to independent when the dependent is moral and the independent is

⁷ For a closely related issue, see Barber's case against explaining the puzzle by appealing to moral supervenience (2013, 643–4).

non-moral—but that bare insistence merely reasserts the thing to be explained. Looking ahead, part of our final diagnosis will indeed appeal to this dependence, but only to explain some of the puzzling cases, and even here we'll explain why it's not the dependence itself that explains why the inferences are problematic.

7.4.2 Moral Science is a Bankrupt *Policy* of Belief Revision

One diagnosis of what's wrong with modifying our non-normative beliefs in light of our normative beliefs might be to contend that though *individual* inferences of this type might be justified, a *policy* permitting these inferences would be an epistemic disaster, and so we should avoid these inferences individually as well. Such a case might begin by noting that inferences from ought to is amount to modifying our picture of the world to fit our evaluations, when arguably our evaluations are more accurate or justified when they are independently informed by our best evidence about what the world is like. As a result, adopting "moral science" as a research program will make one's evaluations less sensitive to evidence, and hence less likely to be accurate. The problem is then multiplied and infects our non-moral beliefs as we continue to infer both moral and non-moral claims from this less accurate set. *Ought to Is* again helps illustrate the problem: inferring that killing an innocent child never maximizes utility from A and B, and blinds us to the tension between act utilitarianism and general prohibitions of act-types (e.g. killing an innocent child), and so we never update (and presumably improve) our moral outlooks in light of the tension.

But this cannot be the whole of the story about what's wrong with inferring from ought to is. First, this diagnosis only explains why there may be something wrong with a general policy or program of moral science—it merely explains why such inferences, though legitimate in isolation, might imperil the accuracy of a class of our beliefs. As such, this diagnosis seems to entail only that moral science is dangerous, but the data to be explained are that we've no grounds for accepting the inferred conclusions even in particular cases. Second, and more importantly, the diagnosis implicitly relies on the very thing we're looking to explain. After all, why isn't it equally plausible to insist that accuracy of our non-moral beliefs is imperiled if we don't treat moral science seriously—why shouldn't we analogously complain that removing these inferences from our epistemic toolbox blinds us to appropriately revising our non-normative beliefs? In short, the reply relies on our bias against revising our non-normative beliefs in light of our normative ones—the very bias we're trying to vindicate.

7.4.3 Inconsistent Grounds

While the premises of inferences from ought to is may be logically consistent, we might think they are epistemically incompatible; no reasonable person would take themselves to have reason to jointly accept them. But what's the source of this incompatibility?

Perhaps while the premises are not inconsistent, their grounds are. For example, one might argue that accepting the first premise requires reasoning about morality in a broadly consequentialist manner, and thus one would determine whether actions are right or wrong using empirical information about their consequences. If this is so, when Smart Alec, accepting the first premise, assesses the second premise, the evidence he took himself to have (e.g. a spontaneous emotional response or an intellectual "insight" that the general principle must be true) wouldn't serve as evidence at all; the inference is valid, but the grounds for one premise vanish when the other is accepted.⁸

This solution sits nicely with our intuition that the premises are in tension. Unfortunately, there is little else to be said in its favor. Alec need not be a committed consequentialist to take seriously his cited reasons for the first premise, nor need he be a committed deontologist to take seriously his cited reasons for the second premise. To deny this is just to insist that the premises or their grounds are incompatible, it is not an explanation. Consider a deontologist that thinks we have strict duties both not to lie and to be compassionate, and takes this to imply that telling the truth can never cause pain. Presumably we think that said deontologist's grounds for these premises are incompatible, but the puzzle is explaining the source or nature of that incompatibility.

7.5 RESOLVING THE PUZZLE: A DILEMMA FOR MORAL SCIENCE

We propose that every moral science inference faces a dilemma. In what follows, we'll walk through how moral scientists come to face this dilemma.

⁸ This might be seen as resolving our challenge to Barber. We can, tentatively, classify our judgments concerning the premises of *Ought to Is* as M-class judgments, but this classification doesn't survive our accepting some consequentialism; once we do so, we are forced to reclassify our judgment that it is always wrong to kill an innocent child as an m-class judgment, not arbitrarily but as a matter of accepting consequentialism. Furthermore, once we recognize this, we see that the argument exhibits just the circularity that Barber posits; a committed consequentialist's evidence for premise B is all the evidence needed to accept the conclusion of *Ought to Is*.

When one wants to infer from a provisionally justified moral claim(s) to an empirical claim:

1. *Either they are committed to treating all the moral premises as necessary or they are not.*

This premise shouldn't be all that controversial, as we're not claiming the moral scientist will be specifically committed to either the premise's necessity, or its contingency; instead we're merely claiming they must deny that it's not neither. As we'll see, each option yields a unique problem for moral science. The type of necessity we have in mind is the moral premise(s)'s independence from contingent non-moral truths. So a truth will be necessary in our sense if it's true "come what may," that is, regardless of how the *particular* non-moral facts turn out, and regardless of which *regularities* hold among these non-moral facts (e.g. the rate of gravitational acceleration, whether humans tend to prefer pie to cake, and the effects of eating only pie on lifespan).

It's worth noting that *any* premise that ascribes a moral property to an actual and particular act/object—for example, *Mary acted wrongly*—will be contingent on particular non-moral facts. After all, Mary's acting wrongly is contingent on Mary's psychological and physical capacities and certainly on her existence. So when a prospective moral scientist's premise ascribes a moral property to an actual particular, or entails such an ascription, the premise is invariably contingent in our sense.

Of course, not all prospective premises of moral science are ascriptions of moral properties to actual particulars. For example, the premises of *Ought to Is* do not ascribe any moral property to any actual act. Instead they assert general relations between properties, i.e. that the property of *maximizing utility* is sufficient for being *not wrong* or *permissible*, and that being an instance of *innocent child killing* is sufficient for being *wrong*. By themselves, such claims entail nothing about the actual world—they merely tell us that if something were to have a particular non-moral property it *would* have certain moral properties.⁹ Even so, such general claims can also be contingent when they depend on regularities that hold among particular non-moral facts. For example, consider again a principle like *eating fish is wrong*. One might regard this principle as contingent on whether fishing causes

⁹ Notice, however, that when such generalizations linking distinct non-moral properties (e.g. being an innocent child killing and maximizing utility) with moral properties (in this case, being wrong and being permissible/not wrong) are conjoined they can entail that law-like regularities hold between the relevant non-moral properties in the actual world (e.g. innocent child killing never maximizes utility). Thus, only non-singleton sets of substantive moral principles can have empirical implications.

environmental degradation or pain. But it's also possible to hold that such a principle is necessary. For example, if one simply (yet oddly) thinks this is just a "bedrock" fundamental moral norm, or, if one thinks that it's wrong because *it's wrong to eat any animal* (assuming fish are necessarily animals).

The first horn concerns contingent moral premises:

2. *If any moral premise is contingent, then the inference is circular or rests on a false assumption.*

In Section 7.3, we were dismissive of Barber's conjecture that all moral science inferences were circular because they, of necessity, invoked m-class judgments, i.e. judgments derived, in part, from *a posteriori* empirical assumptions. For one thing we argued that not all instances of moral science seem to invoke such judgments. However, we do think he was right to claim that when the moral premise rests on some contingent fact(s), the inference exhibits a kind of circularity. In order to make good on this claim, let us illustrate with the following bit of moral science:

Bed Bomb

Having overslept, the moral scientist wakes to the rumbling of hungry unattended young children. Without making any explicit inferences, it seems immediately obvious to him that he should get out of bed. His attempts to rationalize sleeping longer fail; so he forms the judgment *I should get out of bed* and regards it as justified. Never missing a chance to do a little research, our scientist notes the following conditional also seems justified: if a weight-sensitive bomb were attached under my bed, it wouldn't be true that I should get out of bed. The moral scientist then recognizes that these apparently justified claims jointly entail there is no such bomb, he has "discovered" new grounds for believing that there's no such bomb. Feeling tranquil in his increased sense of safety and accomplishment, the moral scientist inadvertently falls back into sleep.

Here, our moral scientist's inference takes the following form:

- I. I should get out of bed.
- II. I shouldn't get out of bed if a weight-sensitive bomb is under my bed.
- III. Therefore, a weight-sensitive bomb is not under my bed.

The inference is valid and uses premises that are ordinarily justified. Nevertheless, the moral scientist's justified beliefs don't support the entailed conclusion. Why not? The answer, we think, is fairly obvious once we consider *how* his first premise could be justified.

First, the moral scientist might be implicitly applying a principle or rule to his circumstances. For example, he might think that he should get up because if he does not get up, he'll break his promise to feed the children before 10.00 a.m., and one should do what will prevent breaking one's promise, or that staying in bed would violate some parental duty, or a duty to not choose minor comforts when doing so puts others at risk, etc.

Whatever particular principle(s), *P*, he uses, and whatever he takes his particular circumstances ($d_1, d_2, d_3 \dots$) to be, his implicit reasoning can be characterized like this:

- i. Under my circumstances ($d_1, d_2, d_3 \dots$), if I get out of bed then I'll comply with *P*.
- ii. When in circumstances like mine ($d_1, d_2, d_3 \dots$), one should comply with *P*.
- iii. Therefore, I should get out of bed.

We have to add (ii) because his principle might be in some way “hedged”—that is, he might assent to a principle (e.g. *one should try to get out of bed as soon as one wakes*) and realize that sometimes one should not comply with the principle/generalization—there are exceptions (e.g. when doing so will detonate a bomb).¹⁰

Now, putting it all together, the moral scientist's pattern of reasoning would look like this:

- i. Under my circumstances ($d_1, d_2, d_3 \dots$), if I get out of bed then I'll comply with *P*.
- ii. When in circumstances like mine ($d_1, d_2, d_3 \dots$), one should comply with *P*.

I. I should get out of bed (i–ii).

II. I shouldn't get out of bed if a weight-sensitive bomb is under my bed.

III. A weight-sensitive bomb is not under my bed (I–II).

We're now in a position to appreciate the problem more precisely. Either ($d_1, d_2, d_3 \dots$) entails there's no bomb under the bed, or it does not. If it does, then the reasoning is objectionably circular, the scientist's moral reasoning plays no role in reaching the conclusion: the inference is valid only because his characterization of the circumstances entails the final conclusion. On the other hand, if ($d_1, d_2, d_3 \dots$) do not entail there's no bomb under the bed, then the premises cannot all be true—the conjunction of (i) and (ii) will conflict with (II). The conjunction entails that premise I is true provided that ($d_1, d_2, d_3 \dots$) and (II) entails that (I) is false when there's a bomb. So, ($d_1, d_2, d_3 \dots$) must either entail that there's no bomb

¹⁰ For those readers who prefer to think of moral principles as a kind of generalization rather than a rule, and hence balk at the use of “comply” in the premises, just substitute claims about compliance or violation with claims about the act type in the principle. For example, if the principle is it's wrong to break a promise, premise (i) would become: i*. Under my circumstances ($d_1, d_2, d_3 \dots$), if I get out of bed then I'll avoid breaking a promise.

or a premise must be false. Thus, the reasoning will always either be circular or employ a false assumption.¹¹

Our example assumes that the moral scientist is implicitly reasoning from a moral principle; will this type of strategy work even if the moral scientist's premise is non-inferentially justified? The issue is crucial because Barber claims the realist can resolve the puzzle only by claiming that, in general, moral verdicts—i.e. claims that ascribe moral properties to actual particulars—can only be justified by deriving them from principles that are not contingent on any empirical assumption. Claiming otherwise, he thinks, opens the door for moral science. However, this isn't so; non-inferentially justified moral verdicts don't open the door to moral science. Even if the moral scientist's premise is justified non-inferentially, the very same sort of problem will arise. To see this, let's return to the *Bed Bomb* case.

First, notice that it's not only plausible, but likely, that the moral scientist's judgment, *I should get out of bed*, can be justified without inferring it from a moral principle or any other moral claim. Suppose that the moral scientist is among the unreflective who don't countenance any moral principle, or among the hyper-reflective who reject moral principles for philosophical reasons, or that he has heretofore made no substantive moral judgments at all. In each case, he can't infer his premise from a moral principle, and in the third case, he can't infer it from any moral claim at all. Nevertheless, it seems to us that his premise may have some justification. Plausibly, if his belief that he should get out of bed were merely sensitive to, but not inferred from, the unsupervised and hungry children, or recalling his promise, then his belief is sensitive to the right considerations—and hence at least somewhat justified. Such sensitivity presumably obtains provided that he either believes the premise in virtue of believing these non-moral facts, or even if he believes it because the facts *seem* or *appear* to him (e.g. he forms the judgment immediately upon hearing the children, or upon replaying his assurance “don't worry, I'll feed them by ten” in his mind). As such, the sensitivity need not involve employing a moral

¹¹ One might worry that our diagnosis won't apply to all reasoning from a single moral principle because not all reasoning from a moral principle involves premises like (ii)—“when in circumstances like mine ($d_1, d_2, d_3 \dots$), one should (or must not) . . .”. Specifically, when the agent takes the principle to be exceptionless and always entails sufficient grounds for doing (or refraining), a premise like (ii) is unnecessary: one should (or must not) in any and all circumstances. Thus, one might suspect that we need a separate diagnosis for why one can't use a principle of this type to do moral science. This suspicion is misplaced for a very simple reason: one cannot reach a descriptive conclusion from a principle of this type. Solitary moral principles have no empirical entailments; without further moral assumptions (including other principles) they're, in isolation, useless for moral science. Indeed, if our moral scientist used a principle of this type he couldn't employ a premise like (II) while holding that his act falls under the scope of the principle.

principle, cognizing these facts as moral reasons, or even inferring the conclusion from non-moral beliefs.

On this picture, the moral scientist's premise "I should get out of bed" can be justified non-inferentially. But does this mean, as Barber suggests, that he could now do legitimate moral science? No. The same problem appears here as when he reasons by applying a principle: the moral scientist would either be committed to a false premise or make a non-moral assumption that already entails the conclusion. To see this, let's remind ourselves of the moral scientist's basic inference:

- I. I should get out of bed.
- II. I shouldn't get out of bed if a weight-sensitive bomb is under my bed.
- III. Therefore, a weight-sensitive bomb is not under my bed.

Here again, we need a story about how (I) is justified. If it is justified non-inferentially, then there are no discrete facts moral or non-moral (such as, the children are hungry or I ought to keep my promise) the scientist reasoned from. Even so, surely the belief, if justified, must somehow be formed in response to, even if not derived from, the scientist's view of the circumstances he took himself to be in. If he denied his belief had anything to do with endangered children, broken promises, the time of day, or even being in bed, then it would be hard to make sense of his belief at all, and even harder to regard it as justified. Instead, in any normal case, one is disposed to say something like "It seemed like I ought to given the circumstances I took myself (or seemed) to be in"—and, again, this certainly could be true even if one didn't infer the conclusion from any particular element(s) of one's apparent circumstances, and even if one can't identify which elements influenced one's conclusion.

But the same problem rears its head again: however the moral scientist construes his circumstances, it either includes or entails the absence of the bomb or not.¹² If it entails it, then the moral scientist's moral premises do no work. If it does not entail it, then from his perspective it's false that he should get out of bed under the apparent circumstances, instead he must have been disposed to hold something like—*I should get out bed under these circumstances only if there's no bomb.*

What's really doing the work in these cases isn't merely the premise's contingency but also the assumption—call it the "mediated access assumption"—that our epistemic access to such premises is ultimately conditional

¹² We'll get different answers depending on how we interpret "takes his circumstances to be," for example depending on whether we include only what seemed to him at the time or admit anything he'd be disposed to accept at the time.

on independent access to what they are contingent on. For example, if an obligation to ϕ depends on one's ability to ϕ , the principle says that any justification for the claim, *I must ϕ* , will be conditional on independent justification for *I'm able to ϕ* —such that evidence or grounds for concluding that *you ought to* provides no reason to believe *you can*, unless it also would be evidence or grounds for concluding that *you can*.¹³ More specifically then, the mediated access assumption claims that when a moral claim, M, is contingent on the truth of some empirical claim(s) E, no grounds for believing M can license believing E unless the grounds independently license believing E. Before defending our use of the principle, a few clarifications are in order.

First, one might suspect that the principle's appeal to "non-moral facts that the moral verdict is *contingent on*" implies a division between the moral and non-moral that's inconsistent with metaethical naturalism. Causal contingency of x on y , for example, does apparently entail that x and y are distinct. We don't have this kind of contingency in mind. What we mean to denote are those features whose truth or falsity bear on the truth of the premise. Thus, if we assume being wrong *just is* being some descriptive feature D , the verdict that x is wrong will clearly be contingent in our sense on whether x is D . Indeed, the mediated access principle is not in tension with metaethical naturalism, metaethical naturalism entails it.

Second, note that the relevant non-moral features need not refer to right- or wrong-making features of the act. A moral verdict's truth can apparently turn on facts that are not right or wrong "makers." For example, the moral scientist presumably denies that absence of a bomb explains why he should get up, though by his lights, were there one then he shouldn't get out of bed.

Third and finally, we concede that perhaps a moral belief could be justified by testimony or induction, and that while these types of justification would require access to some non-moral facts (e.g. hearing the testimony and being

¹³ "Conditionally justified" here can be interpreted two ways: one meaning one can't be justified at all in the ordinary sense without some justification for the non-moral facts the claim depends on, or alternatively, meaning that one may be justified, but only in a conditional sense—i.e. the belief enjoys some variety of epistemic merit, but not one that further licenses drawing unqualified inferences from that belief. Something along the second interpretation seems most plausible. After all, if one realizes something ought to be done, provided they can, one seems—all else being equal—entitled to believe and act as if they ought to, but one shouldn't then use this belief to modify her beliefs regarding, for example, what she can physically do, or her beliefs about the act's likely consequences. So, in the *Bed Bomb* case, we happily acknowledge that perhaps the moral scientist knows *he ought to get up*, and that *there is no bomb* (if there is no bomb), or that perhaps both of these beliefs are justified (whether there is a bomb or not). What we reject is that he's *entitled to infer* there's no bomb from I ought to get up, or treat it as evidence that he ought to get up.

aware the present case is non-morally similar to some previous cases) these need not be facts on which the moral claim is contingent. This concession is not in conflict with the principle; the principle asserts that *our* epistemic access to such premises is ultimately conditional on independent access to what they are contingent on. So, yes, perhaps reliable testimony that *I ought* may justify my inference that *I can*, but for this to be possible, someone, at least, must have had independent access to facts about my capacities.

Denying the mediated access principle entails that one *could* get direct evidence for the moral premise and hence reasons to accept the non-moral facts on which it is contingent; the fact that we find Barber's *Permissible to Painless* and our *Bed Bomb* examples ridiculous reveals our commitment to the principle. But won't a complete answer to the puzzle defend the mediated access assumption? No. We can claim that our acceptance of the principle is central to resolving our puzzle without defending it. But, yes, it would be nice to *vindicate* our rejection of moral science, not just explain it. And there are putative proofs of the principle.¹⁴ But a proof seems dialectically unnecessary; even friends of moral perception implicitly accept some version of this thesis. Duty, like beauty, may be seen directly—but seeing either presumably requires access to empirical features. One can't "main-line" a painting's beauty without being sensitive to *any* facts regarding its color, shape, composition, subject, aims, historical context, artist, contrasts (etc.); and so it is with an act's moral status and the non-moral features it's thought to be contingent on. But before moving on to the next horn, allow us to gesture further at why denying the principle seems absurd.

Even if a case against the mediated access assumption could be made, it would not vindicate any folk notion of a moral sense, instead it would be an incredible revision. To be sure, we often claim to "just see" that certain options are forbidden or required—spontaneously, non-inferentially, and without explicit rationale. For all we know, these judgments may be reliable, but they are not plausibly interpreted as claims to have detected a moral fact independent of its non-moral truth conditions. When I stand on the pond's shore and "just see" I must save the child, I'm not acquiring independent evidence that I won't get stuck in quicksand en route, or that what I see is, indeed, a child drowning. Our sense, if we have one, would operate *under* these assumptions—our view of the circumstances—rather than providing

¹⁴ See, for example, Zangwill (2006), who argues it's a conceptual truth that ultimately rests on the necessity of the truth that moral properties obtain in virtue of their descriptive properties, Faraci (2015), who purports to prove any defensible view of moral perception obeys the principle, and for the most famous rhetoric in defense of the thesis see Dworkin (1996, 104–6).

evidence *for* them. Furthermore, moral judgments are typically made in practical contexts—contexts where the question is *what to do?* Shall I jump in the pond? Here one is trying to determine what *would or will be* right, permissible, or wrong. The point is that *acts* have moral properties, and when we face practical questions, nothing has been done yet—there’s nothing to instantiate the property. In these contexts, if our moral sense really worked by unmediated or unconditional detection of a moral fact, then, incredibly, events in alternative futures would explain the current outputs of my moral sense. While these concerns don’t prove our assumption, we hope it shows where the burden lies. Indeed, if the assumption were false, moral science, particularly as employed by the moral scientist in *Bed Bomb*, would be faultless. So because we endeavor to explain why moral science is problematic, it’s dialectically reasonable to assume the mediated access assumption is true—otherwise there would be nothing to explain.

As it turns out, our treatment of the dilemma’s first horn is much like Barber’s “realist-friendly” solution. In fact, Barber’s solution works like ours because his “strategic division” entails mediated access: if moral verdicts can enjoy justification only via deduction from *a priori* principles and empirical application conditions, then their justification depends on independent access to those empirical conditions. Crucially, however, the mediated access principle does not entail the strategic division: it is weaker and more ecumenical. For example, it is consistent with non-inferential, perceptual, and case-based moral epistemologies and a particularist moral metaphysics; the strategic division is not. But the biggest difference, of course, is that on our view appeals to circularity are only a *partial* solution—they cover only some of the problematic yet difficult to diagnose inferences from *is* to *ought*. In Section 7.3 and earlier in this section we saw that not every moral science inference invokes an *m*-class judgment and similarly not every inference rests on some moral premise that, by the moral scientist’s own lights, is contingent. It’s these inferences, inferences like Alec’s, that inspired our interest in this topic. The second horn of our dilemma covers this branch of moral science.

3. *If they are committed to treating all the moral premises as necessary, they are forced to deny the contingency of a claim known to be contingent.*

In some cases, like Alec’s, all the moral premises of the inference from *Ought to Is* are taken to be necessary truths. Arguably, only this class of moral science is worthy of the name. The inference from moral to non-moral in other instances is, given the argument in this section, illusory; the inference can be justified only when and because the moral scientist may non-morally infer the same conclusion. What can we say to the true moral scientists like Alec? Our response, put simply, is that the moral scientist has substantial

grounds for denying the conclusions of such inferences and so has grounds to reject at least one of the premises.

The grounds for rejecting the conclusions of moral science is based on the grounds one has for denying a logically stronger conclusion that follows from the premises of such inferences where the premises are taken as necessary. Consider a logical consequence of taking the premises of *Ought to Is* as necessary. Consider *Ought to Must*:

Ought to Must

It is never wrong to do what maximizes utility.

It is always wrong to kill an innocent child.

C*. Therefore, necessarily, killing an innocent child will never maximize utility.

The premises of *Ought to Must* are the same as the premises of *Ought to Is*, the difference is that the conclusion, call it MUST, is explicitly not about the actual world but about all possible worlds. Recognizing that the premises of *Ought to Is* are taken as necessary to yield MUST helps diagnose the puzzle. In short, we know that MUST is false, and knowing that MUST is false it follows that at least one of the premises is false.

To avail ourselves of this solution is to avail ourselves of a bit of modal knowledge, but neither our claimed knowledge nor its modal status should give us pause. Let's start with the fact that the claimed knowledge is modal. First, it is the sort of modal knowledge to which we are, we think, typically and often entitled. Indeed, our knowledge that it is possible that killing an innocent child could maximize utility seems as secure as our knowledge that the human species might never have evolved, that the current climate crisis might have been prevented, or that we might not have written this chapter. These all seem like ordinary modal propositions we take ourselves to know.¹⁵ And while we might be rightly skeptical about modal knowledge concerning, for example, the possibility of a perfect being (Van Inwagen 1998) or the possibility of beings that are identical to us in every way except that they lack consciousness (Chalmers 1997), the possibility of maximizing utility by killing an innocent child doesn't seem to be of this type—it doesn't

¹⁵ In what follows we mostly adopt a form of non-skeptical modal rationalism, a modal epistemology on which modal knowledge is *a priori* and on which conceivability (Yablo 1993; Chalmers 2002) or imagination, or consideration of counterfactuals (Williamson 2008) are sources of moral knowledge. We adopt such a view, to the extent that we do, only because it comes naturally to us as non-experts in modal epistemology. However, to the best of our knowledge, there is nothing about the modal knowledge we take ourselves to have that is incompatible with or even particularly controversial for those that deny modal rationalism in favor of alternatives (Jenkins 2010; Shalkowski and Bueno 2000; Roca-Royes 2010; Bueno and Shalkowski 2015).

involve a deep insight about metaphysical possibility, but rather an ordinary judgment about physical possibilities.

Another source of skepticism about modal knowledge concerns *a posteriori* necessities.¹⁶ If a proposition is an *a posteriori* necessity, then many of the tools we might use to investigate possibility—for example, conceivability or imagination—are impotent. But the proposition that necessarily it never maximizes utility to kill an innocent child, can't, by Alec's own lights, be of this sort; in this case *MUST* is derivable from propositions Alec takes himself to know *a priori*.

Furthermore, as detailed below, our knowledge that *MUST* is false is at least as secure as our knowledge that "actually, killing a child never maximizes utility" is false; after all, *MUST* entails this claim about the actual world so the probability that *MUST* is true is greater than or equal to whatever probability we assign the claim about our world. So, anyone who takes themselves to be sure that, in our actual world, killing an innocent will sometimes maximize utility, should also take themselves to have knowledge that *MUST* is false.

Finally, notice how one of the ways we might be inclined to respond to Alec presupposes knowledge that *MUST* is false. Recall that in our earlier discussion of Alec we had ourselves insisting that Alec should rethink one of the premises of *Ought to Is* because of his assent to a thought experiment where, by hypothesis, killing an innocent child would maximize utility. We had Alec respond that, yes, he agreed that it was possible that killing an innocent child would maximize utility, but not that it ever actually would. Recognizing that the premises of *Ought to Is* entail *MUST* and that *MUST* is false helps us to explain why a thought experiment really does have argumentative force, and why Alec really must reconsider at least one of the premises (or retreat to taking one or both of the premises to be contingent). If we do not presuppose knowledge that *MUST* is false, it isn't merely that our proposed thought experiment fails to undermine Alec's reasoning, we couldn't even make sense of why we raised it in the first place.

Nor does the solution presuppose any controversial claims about *knowledge*; the solution can be explained instead by talking about our relative confidence that *MUST* is false. Explaining the solution in these terms is a bit more technical, but it has the same force. Again, in short, our credence in *MUST* provides an upper limit on how confident we can be or become in the conjunction of the premises of *Ought to Is*. And, so as long as we are fairly confident that *MUST* is false, we can't increase our confidence in these premises in a meaningful way.

¹⁶ See Kripke (1982) for a discussion of *a priori* necessity. See Yablo (1993) for a discussion of how such necessities undermine modal knowledge.

To see why, it's useful to consider how our credences in the premises of *Ought to Must* relate to our credence in MUST itself. While evidence isn't transitive over entailment generally, there are circumstances where evidence is transitive (Roche and Shogenji 2013). For example, assume you are equally confident in two propositions about a (normal) playing card that is face down in front of you: (1) the card is a diamond; (2) the card is red. Since (1) entails (2) one can't rationally be more confident in (1) than (2). But then any evidence we take ourselves to have that would increase our confidence that the card is a diamond must also transfer, or serve as evidence that the card is red. This explains why our credence in MUST sets an upper limit on how confident we can become in the conjunction of the premises of *Ought to Must* (and so the premises of *Ought to Is* taken as necessary truths). Since the premises entail MUST, we can never be more confident in the conjunction of the premises than we are in MUST.¹⁷

Now that the relationship between MUST and the premises is clear, we can see how MUST suppresses our confidence in the premises; assigning a low credence to MUST means assigning a similar or lower credence to A and B. We need not adopt a view about whether we *know* MUST to be false to criticize our inference, we must only admit what is even more obvious; that we are very skeptical of MUST.

Of course, Alec might concede that his credence in MUST is indeed low and that his credence, therefore, in the conjunction of the premises of *Ought to Must* is also low but argue that his reasoning about the premises has led him to increase his confidence in MUST slightly rather than to discount his evidence for the premises. What are we to say to Alec who insists on realigning his credences in this manner?

So long as Alec's credence in MUST is significantly lower than a half, it isn't clear that we have to say anything to Alec. Alec was pretty sure that MUST was false and he remains pretty sure that MUST is false even if a little less so. He certainly doesn't have grounds to believe MUST or even to be agnostic about it.

This horn of the dilemma has been couched in terms of a particular moral science inference, *Ought to Is*, but we hope it is clear why it applies more generally. Whenever all the moral premises of a moral science inference are taken as necessary, there will also be an entailment that is necessary. But, moral science inferences are troubling precisely because they seem to

¹⁷ If one is certain that MUST is false, i.e. assigns MUST a credence of 0, then it is impossible to gain evidence for the conjunction of the premises at all. Probabilities of 0 and 1 are "sticky" in that no amount of evidence can move us from that assignment (Howson and Urbach 2006; Sober 2008, chapter 1). But then we must also assign the conjunction of the premises a credence of 0.

provide information about matters that are contingent on empirical facts. Every moral science inference with premises taken to be necessary will entail the necessity of something we take to be contingent.¹⁸

7.6 CONCLUSION

It seems that no matter what our moral scientist takes the modal character of his moral premises to be, it has such a character; and whether they're jointly contingent or necessary, he is in a position to see his project will fail. Of course, the argument can't save every aspiring moral scientist. One can, after all, deny the mediated access principle, or antecedently possess very strange modal views. But if that's what it takes to be rational and do moral science, then we can finally explain why it's not for us; and do so without positing collective delusion or rejecting major metaethical positions. Perhaps best of all, we can finally explain to Alec why he needs to revise one of those premises!¹⁹

¹⁸ Jamie Dreier, in conversation, has wondered about the possibility of moral science inferences that invoke only a single premise. While a great many individual moral premises have descriptive entailments, for example, any single premise that attributes a "thick" moral predicate to an actual object or single premises such as "I ought" may entail "I can"; in these examples it is clear that any descriptive entailment may be inferred only when whatever justifies the premise would also justify it. And that's just what the mediated access principle predicts. But the more tantalizing kind of case that Jamie had in mind was of a single-premise argument using a general moral principle. It is harder to see how one could construct a single-premise argument using general moral principles (as opposed to particular moral judgments) since these principles don't, by themselves, entail anything about the way the world is because they don't tell us that any moral properties are actually instantiated. It is only by combining two or more moral principles that we can derive conclusions about what must be, since only then can we derive restrictions on what can actually be instantiated. So while some have argued that Utilitarianism entails that the universe is finite in duration (see Nelson 1991), and that this looks like a one premise inference with a putatively necessary premise, the inference will be invalid unless we also assume that some act actually is or actually will be wrong/right, or valid only if we claim Utilitarianism itself is a conjunction of a monistic principle and the claim that the principle actually applies to some act. So understood, the inference would fall prey to the first horn: one's justification for thinking some actual act was right/wrong or thinking that the principle applied would require independent evidence that the universe is finite in duration.

¹⁹ We would like to thank the following for helpful input on this chapter: Ralph Bader, Matt Barker, Marty Barrett, Selim Berker, Matthew Brown, Agnes Callard, Richard Yetter Chappell, David Copp, Dale Dorsey, Jamie Dreier, David Faraci, Patrick Forber, Paul Gowder, Zoe Johnson-King, David Killoren, Matt Kopec, Don Loeb, Pete Murray, Steve Nathanson, Doug Portmore, Bill Roche, Richard Rowland, Ron Sandler, Joshua Schechter, Andrew Sepielli, Russ Shafer-Landau, Scott Simmons, Neil Sinhababu, Rory Smead, Michael Smith, Jacob Sparks, Daniel Star, Michael Titelbaum, Teemu Toppinen, Ralph Wedgwood, Jack Woods, Alex Worsnip, and two reviewers at Oxford University Press. We were fortunate to present this work at the annual Chapel Hill Metaethics Workshop and are grateful for the constructive feedback received there. We'd

References

- Barber, Alex. 2013. "Science's Immunity to Moral Refutation." *Australasian Journal of Philosophy* 91 (4): 633–53.
- Bueno, Otávio, and Scott A. Shalkowski. 2015. "Modalism and Theoretical Virtues: Toward an Epistemology of Modality." *Philosophical Studies* 172 (3): 671–89.
- Chalmers, David. 1997. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Chalmers, David. 2002. "Does Conceivability Entail Possibility?" In *Conceivability and Possibility*, edited by Tamar Gendler and John Hawthorne, 145–200. Oxford University Press.
- Crawford, Jarret T., and Anuschka Bhatia. 2012. "Birther Nation: Political Conservatism is Associated with Explicit and Implicit Beliefs that President Barack Obama is Foreign." *Analyses of Social Issues and Public Policy* 12 (1): 364–76.
- Dorr, Cian. 2002. "Non-cognitivism and Wishful Thinking." *Noûs* 36: 97–103.
- Dworkin, Ronald. 1996. "Objectivity and Truth: You'd Better Believe it." *Philosophy & Public Affairs* 25 (2): 87–139.
- Faraci, David. 2015. "A Hard Look at Moral Perception." *Philosophical Studies* 172 (8): 2055–72.
- Geach, Peter T. 1965. "Assertion." *The Philosophical Review* 74 (4): 449–65.
- Howson, Colin, and Peter Urbach. 2006. *Scientific Reasoning: The Bayesian Approach*. Open Court Publishing.
- Jenkins, C. S. 2010. "Concepts, Experience and Modal Knowledge." *Philosophical Perspectives* 24 (1): 255–79.
- Jones, Jeffery M. 2010. "Conservatives' Doubts about Global Warming Grow." Gallup. <<http://www.gallup.com/poll/126563/conservatives-doubts-global-warming-grow.aspx>>.
- Kahan, Dan M., and Donald Braman. 2005. "Cultural Cognition and Public Policy." SSRN Scholarly Paper. Social Science Research Network. <<http://papers.ssrn.com/abstract=746508>>.
- Kripke, Saul. 1982. *Naming and Necessity*. Harvard University Press.
- McCright, Aaron M., and Riley E. Dunlap. 2011. "Cool Dudes: The Denial of Climate Change among Conservative White Males in the United States." *Global Environmental Change* 21 (4): 1163–72.
- Nelson, Mark T. 1991. "Utilitarian Eschatology." *American Philosophical Quarterly* 28 (4): 339–47.
- Roca-Royes, Sonia. 2010. "Modal Epistemology, Modal Concepts and the Integration Challenge." *Dialectica* 64 (3): 335–61.
- Roche, William, and Tomoji Shogenji. 2013. "Confirmation, Transitivity, and Moore: The Screening-Off Approach." *Philosophical Studies* 168 (3): 797–817.

also like to thank Pea Soup for hosting an initial discussion of the puzzle raised in this chapter. Certainly we have failed to thank everyone who helped, we apologize for any exclusions.

- Shalkowski, Scott A., and Otávio Bueno. 2000. "A Plea for a Modal Realist Epistemology." *Acta Analytica* 15 (24): 175–93.
- Sober, Elliott. 2008. *Evidence and Evolution: The Logic Behind the Science*. Cambridge University Press.
- Van Inwagen, Peter. 1998. "Modal Epistemology." *Philosophical Studies* 92: 67–84.
- Williamson, Timothy. 2008. *The Philosophy of Philosophy*. John Wiley & Sons.
- Yablo, Stephen. 1993. "Is Conceivability a Guide to Possibility?" *Philosophy and Phenomenological Research* 53 (1): 1–42.
- Zangwill, Nick. 2006. "Moral Epistemology and the Because Constraint." In *Contemporary Debates in Moral Theory*, edited by James Lawrence Dreier, 263–81. Blackwell.