

# MORAL SKEPTICISM

Matt Lutz and Jacob Ross

Most people think they know right from wrong. While there may be certain moral gray areas where there is room for reasonable disagreement, we generally think we know many moral truths, e.g., that murder, torture and rape are morally wrong, and that helping the needy is morally good. But are these claims to moral knowledge defensible?

The moral skeptic denies that we can have moral knowledge, that is, knowledge of moral truths. In this paper, we will look at a number of influential arguments for moral skepticism. We will first argue that many of these arguments can be helpfully understood as expressing the same core concern, which we will call the *Connection Concern*. In Section 2, we present a refined version of the Connection Concern. This argument – the Explanatory Trilemma Argument – makes a strong case for moral skepticism. And in Section 3 we examine potential avenues for objection to the argument.

A couple of notes before proceeding. First, we will be assuming that there are such things as moral beliefs, about which it makes sense to ask whether they are true or false, whether they are justified or unjustified, and whether they constitute or fail to constitute knowledge. Second, our focus will be on moral skepticism as a claim about moral *knowledge*. Skepticism about the possibility of moral truth will not be our focus; we will be concerned with arguments that allow, *arguendo*, that there are such truths, but which deny that we are in a position to know them. And other epistemic categories, such as justification, will be relevant only insofar as they are necessary conditions on knowledge. Third, we will be concerned with skepticism about *moral* knowledge, and, in particular, knowledge of *right and wrong*. However, many of the arguments we consider could be extended to other kinds of ethical knowledge, such as knowledge of good and bad or virtue and vice. Fourth, our focus will be on moral skepticism understood as the view that *we can have no moral knowledge*. We will not be asking,

for example, whether there could be an omniscient being who possesses moral knowledge. Instead, we will be asking whether moral knowledge is attainable by us, in our actual circumstances. And fifth, we are concerned with arguments for the unattainability of moral knowledge where this conclusion does not follow *a fortiori* from global skepticism. An interesting argument for moral skepticism will show why *moral* knowledge is particularly troublesome.

## **1. Varieties of Skeptical Argument and the Inescapability of Explanatory Connections**

There are many different arguments for moral skepticism. In this section, we will survey some of the most prominent such arguments, and show that all these arguments can be understood as sharing a core concern, which we will call the *Connection Concern*. While there may be other skeptical arguments that cannot be so understood (for instance, Sinnott-Armstrong's regress argument (Sinnott-Armstrong 2006, Chapter 4) is not a version of the Connection Concern), the Connection Concern is so central to such a large number of skeptical arguments that it will be our main focus in this paper. In its most schematic form, the Connection Concern argument can be represented as follows:

(P1) S knows that P only if S's belief that P is connected to the fact that P in the Right Kind of Way.

(P2) For any moral proposition P, it is not the case that our belief that P is connected to the fact that P in the Right Kind of Way.

(C) Therefore, for any moral proposition, P, it is not the case that we know that P.

We will set aside a critical discussion of the Connection Concern argument until Section 2. For now, our interest is in the relationship between this argument and other varieties of skeptical argument.

### **1.1 Moral Disagreement**

By far the most popular and influential argument for moral skepticism is the argument from moral disagreement. There is a substantial amount of disagreement on what counts as right or wrong between cultures and over times. On this basis, some have concluded that moral knowledge is impossible.

Different people seem to have opposing moral beliefs, and who is to say which moral beliefs are

correct? Yet disagreement between individuals, or between members of different cultures, over some matter of fact does not immediately preclude knowledge of that fact (see Shafer-Landau 2004, Ch. 1; Locke (this volume)). If Alice and Brian disagree about whether human beings evolved from non-human animals, that does not imply that there can be no knowledge of human evolution, nor does it imply that neither Alice nor Brian possesses such knowledge. If Alice, who is aware of the scientific evidence supporting human evolution, believes that humans evolved from other animals on the basis of this evidence, then she may know that humans so evolved in spite of the fact that Brian, who is unaware of this evidence, believes otherwise.

Thus, if the skeptic wants to argue that moral disagreement leads to a skeptical conclusion about morality, some effort must be made to show why disagreement about morality is particularly worrisome. So we must ask: when, in general, does disagreement create epistemic problems? Often, when we learn that someone disagrees with us about some proposition, it will be rational for us to become less confident that this proposition is true. In other words, disagreement can often act as a *defeater*. Defeaters are often thought of as falling into two (possibly overlapping) kinds: *rebutting defeaters* and *undercutting defeaters* (Pollock 1987). Rebutting defeaters are defeaters that operate by providing evidence for a contrary proposition. Undercutting defeaters, on the other hand, attack the evidential connection between our old evidence and the proposition in question. They provide new evidence in light of which the old evidence no longer provides (as much) support for the belief that P. Testimony that not-P can provide a defeater for one's belief that P of either of these two kinds. Suppose, for example, that Albert believes that Betty will be at the party because Carl told him so. But then Diane comes along and tells Albert that Betty will not be at the party. Normally, Diane's testimony will provide evidence that Betty will not be at the party, and so it will be a rebutting defeater for his belief that she will be at the party. But suppose that Albert knows that Diane's formed her belief on the basis of Carl's testimony to her. In this version of the case, Diane's assertion that Betty will be at the party indicates that Carl has been saying conflicting things to different people, and so it indicates

that his testimony is unreliable. In this way, Diane's testimony provides an undermining defeater for Albert's initial belief.

Moral disagreement might be thought of as providing a defeater for our moral beliefs of either of these kinds. However, if we adopt that view that moral disagreement merely provides a rebutting defeater for our moral beliefs, then we should not regard such disagreement as supporting moral skepticism. For on this view of disagreement, someone's testimony that not-p serves as a defeater for our belief that p by providing evidence for not-p. And so, on this view, moral testimony provides evidence for the moral claim asserted. But if this is right, then presumably, other things being equal, the more people sincerely assert a given moral claim, the more reason we will have to believe it. Hence, if there are many more who agree with you on a question of ethics than disagree, then the defeat that comes from the few dissenters will be swamped by the support provided by the much larger number who agree with you. Consequently, while the contrary evidence provided by disagreement may jeopardize our justification for our more controversial moral beliefs (e.g., our beliefs about abortion or affirmative action), such contrary evidence won't jeopardize our justification for our relatively uncontroversial moral beliefs (e.g., our beliefs about the wrongness of torturing babies). And so, if moral disagreement provides only a rebutting defeater, moral disagreement will not support a broad moral skepticism.

In order for moral disagreement to support general moral skepticism, disagreement concerning *controversial* topics would need to cast doubt on our moral beliefs concerning *uncontroversial* topics. Thus, for example, disagreement about affirmative action would need to cast doubt on our beliefs about torturing babies. And it may do this by providing an *undercutting* defeater. For even if I don't have a good understanding of how my moral beliefs are formed, I may have reason to suppose that my moral beliefs and the moral beliefs of others are formed on the basis of similar mechanisms. Hence, learning that others disagree with me about a range of moral issues may call into question the reliability of the mechanisms underlying my moral beliefs, and hence call into question the accuracy of all the beliefs,

including the uncontroversial ones, formed on the basis of these mechanisms. And if our moral belief-forming mechanisms are unreliable, then it would seem that our moral beliefs are not connected, in the right sort of way, to the moral facts.

To sum up, if the Argument from Disagreement is to be successful in supporting a general skeptical conclusion, then moral disagreement must be thought of as providing an undercutting defeater for our moral beliefs, and hence as indicating that our moral beliefs are not connected, in the right sort of way, to the moral facts. Hence, to the extent that the Argument from Disagreement succeeds, it does so by raising the Connection Concern.

## **1.2 Divergence in Mature Theory**

Another common skeptical argument proceeds from the claim that there has been little or no convergence over time between competing moral theories. The differences between consequentialism and deontology, for instance, and the arguments for and against each, have long been fairly well understood. Yet there has been no convergence in expert opinion as to which of these two theories is superior, and there are no signs that such a convergence will be forthcoming any time soon. This makes moral theory notably different from scientific theory, where convergence and progress are *de rigueur*. (This concern is plausibly the underlying motivation for Derek Parfit's attempt to reconcile various normative theories (Parfit 2013).)

While this is a disconcerting feature of moral discourse, the lack of convergence between mature moral theories is simply an instance of protracted moral disagreement among moral theorists. And so we should understand the Divergence in Moral Theory Argument as a special case of the Moral Disagreement Argument. As such, insofar as it is successful, it is also just a special case of the Connection Concern.

In fact, it is useful to reflect on the Divergence in Moral Theory Argument because it is a particularly good illustration of why concerns about disagreement point to the Connection Concern.

Prolonged, intractable disagreement about moral matters persists, even among expert theorists. What could explain this? One possibility is that there are no moral facts to begin with, so of course the experts are spinning their wheels. If there are moral facts, on the other hand, one likely explanation of this continued disagreement between experts is the hypothesis that even experts are not forming beliefs in a manner that is conducive to finding the truth of the matter (although this is not the only explanation – see Enoch 2012, Ch. 8; Huemer 2005, Ch. 6). And so such disagreement suggests that, even under the best epistemic conditions we can hope for, our beliefs aren't appropriately connected to the moral facts.

### **1.3 Evolutionary Debunking and Off-Track Influences**

Another argument for moral skepticism that has gained substantial popularity in recent years is the Evolutionary Debunking Argument (Street 2006; Joyce 2006). According to the Evolutionary Debunking Argument, we know what explains our moral beliefs: evolutionary pressures. We have the moral beliefs we do because having dispositions to form such moral beliefs was conducive to the survival of our ancestors. Because our moral beliefs track evolutionary pressures so neatly, it would be a massive coincidence if those beliefs were *also* reliably correlated with a mind-independent moral truth.

The Evolutionary Debunking Argument has taken a number of different forms in recent years. Street and Joyce present the debunking concern in different ways, and draw different conclusions from their arguments (Joyce thinks the argument supports moral error theory; Street deploys it in favor of constructivism). But while there are many different ways to frame the debunking argument, we see it as another instance of the Connection Concern. Whereas arguments from disagreement give us a reason to suspect that our moral beliefs don't reliably track the moral facts, the Evolutionary Debunking Argument looks at what our moral beliefs *do track*, and suggests that what is being tracked has nothing to do with morality. Our moral beliefs are not connected to the moral truth, because they are connected to something else, i.e. facts about which kinds of moral dispositions happen to have promoted the survival of our Paleolithic ancestors.

In a similar vein, Walter Sinnott-Armstrong has argued that our moral beliefs do not track the moral truth, because they track other kinds of psychological pressures (Sinnott-Armstrong 2006: Ch. 9). Drawing on much contemporary psychological research, Sinnott-Armstrong shows how our moral beliefs can be affected by a huge number of morally irrelevant factors. For instance, we are somewhat more likely to judge that an action is morally wrong when we are in a dirty environment than a clean one, because dirty environments evoke emotions of disgust and those emotions affect our moral reasoning. We are also subject to “framing effects” when making our moral judgments: we are liable to describe the same action as morally wrong or morally neutral depending on whether the action is described as a loss from a high benchmark or a gain from a low benchmark.

Sinnott-Armstrong musters a battery of psychological experiments that demonstrate the many ways in which our moral beliefs can be influenced by factors that are irrelevant from the standpoint of morality: the moral permissibility of a given action may depend on many things, but it does not depend on whether there is rotting food in your vicinity when you hear about the action. And this is, once again, an instance of the Connection Concern. What these experimental results suggest is that our moral beliefs are not connected to the moral facts in the right kind of way, since, in forming them, we are at the mercy of irrelevant environmental factors.

We do not claim to have shown that all of these skeptical arguments *must* be interpreted as special cases of the Connection Concern. But the Connection Concern is central to so many different arguments for moral skepticism that it is worth isolating and examining in its own right, and attempting to present its strongest formulation. The success or failure of such an argument will tell us much about the prospects for moral skepticism. So our task in the remainder of this paper will be to identify and evaluate the strongest version of the Connection Concern.

## **2. A Refined Version of the Connection Concern Argument**

Recall that the initial version of the Connection Concern Argument we presented ran as follows:

(P1) S knows that P only if S's belief that P is connected to the fact that P in the Right Kind of Way.

(P2) For any moral proposition P, it is not the case that our belief that P is connected to the fact that P in the Right Kind of Way.

(C) Therefore, for any moral proposition, P, it is not the case that we know that P.

In this section, our aim will be to move from this schematic outline to a clear and defensible argument.

## 2.1. Defining the Relevant Connection

In order to refine this argument, we will first need to specify what we mean by the "Right Kind of Way." To do so, we will need to find some kind of connection between belief and the facts which is plausibly required for knowledge, and which is arguably absent in the case of moral beliefs.

To find such a connection, it may be useful to consider Gettier cases in which someone appears to have justified true belief that fail to constitute knowledge, and where what explains why these beliefs do not constitute knowledge is that they fail to be connected to the corresponding facts in the right kind of way. Consider the following cases:

**Coin Conundrum:** Your boss tells you Jones will be getting a promotion. And Jones tells you that he has ten coins in his pocket. On this basis, you infer that the person who will receive the promotion has ten coins in his pocket. While you have not counted the coins in your own pocket, it happens that they are ten in number. And it turns out that it is you, and not Jones, who receive the promotion. And so your belief is true (Gettier 1963).

**Sheep Shenanigans:** You are standing in a field looking at a dog which has been cleverly disguised as a sheep. On this basis, you form the belief that there is a sheep in the field. As it happens, elsewhere in the field there is indeed a sheep that is hidden from your view. And so your belief is true (Chisholm 1966).

In both these cases, it seems that, while you have justified true belief, it does not amount to knowledge.



One natural thought is that your belief isn't knowledge because you were *lucky* in arriving at a true belief. But this can't be quite right. If you just happen to turn your head in the right direction just as a meteor appears in the sky, then you are lucky in arriving at the true belief that there is meteor in the sky, but your belief amounts to knowledge all the same. What distinguishes the meteor case from the other two cases just considered is that, while, in all three cases, you are lucky to form a true belief, in the meteor case is *no mere coincidence* that your belief is true. After all, the presence of the meteor caused your visual experience, which caused your beliefs. Hence, you believe that there is a meteor in the sky *because there is a meteor in the sky*. By contrast, in the other two cases, it is a mere coincidence that your belief is true. For, in the coin case, you believe that the person who will receive the promotion has ten coins in his pocket not because this person really does have ten coins in his pocket, but because someone else has ten coins in his pocket. Similarly, in the sheep case, you believe that there is a sheep in the field not because there is a sheep in the field, but because there is a dog disguised as a sheep in the field.

These cases suggest that, in order for your belief to count as knowledge, it can't be a mere coincidence that this belief is true. We should note that this is not meant to be a sufficient condition for knowledge, or even to yield a sufficient condition when added to justified true belief. However, it does appear to be plausible as a necessary condition for knowledge. Moreover, it should be widely accepted as such, since many accounts of knowledge that have been offered in the literature entail it. For many such accounts require that, in order for one's belief that *p* to count as knowledge, there must be some kind of counterfactually robust connection between one's belief that *p* and the fact that *p*, so that they covary, in some specified way, across possible worlds (Particularly prevalent here are "sensitivity" accounts – see Nozick (1981) – and "safety" accounts – see Sosa (1999)). And since mere coincidences tend not to be counterfactually robust, these accounts will typically imply that, in order for one's belief to count as knowledge, it can't be a mere coincidence that one's belief is true.

*A mere coincidence* is a concurrence without an explanatory connection. Thus, for example, it is

a mere coincidence that Thomas Jefferson and John Adams died on the same day. If, however, for example, the death of one caused the death of the other (e.g., because the second died of grief upon hearing the news of the other's death) or if both deaths had a common cause (e.g., because they were both killed by the same explosion), then it would not be a mere coincidence that they both died on the same day. Thus, the concurrence of A and B will constitute a mere coincidence just in case it is not the case that A explains B or that B explains A or that some other thing explains both A and B.

Consequently to say that, in order for one's belief to count as knowledge, the concurrence between one's belief and the truth can't be a mere coincidence is equivalent to the following:

**Explanatory Connection Condition:** S knows that p only if either (i) the fact that p explains the fact that S believes that p or (ii) the fact that S believes that p explains the fact that p or (iii) some further fact explains both the fact that p and the fact that S believes that p.

This will be the condition on knowledge that interests us in what follows.

## 2.2. Pure and Impure Moral Propositions

Assuming the Explanatory Connection Principle, all we'd need in order to derive moral skepticism is the following:

(2\*) For any moral proposition p, it is not the case that our belief that p is explained by the fact that p, or vice versa, and neither is it the case that the fact that p and our belief that p share a common explanation.

Unfortunately for the skeptic, however, this assumption is highly questionable. Consider, for example, our belief that Pol Pot acted wrongly. Arguably, we have this belief because Pol Pot acted wrongly. After all, if Pol Pot had minded his own business and avoided all wrongdoing, we wouldn't now believe he had acted wrongly. And so it seems that there is indeed an explanatory connection between the fact that Pol Pot acted wrongly and our belief that he so acted (see Sturgeon 1988). While the existence of such explanatory connections is controversial, it would be better for the skeptic to make an argument that does not depend on denying them.

One way to solve this problem would be to identify a subset of moral propositions that plausibly do not stand in such explanatory connections with our beliefs. In order to do so, it will be useful to introduce some terminology. Let us define a *complete moral theory* as a theory that specifies, for every maximally specific consistent descriptive characterization of a world, how things would stand morally in a world satisfying that description. And let us define a *pure moral proposition* as a disjunction of complete moral theories. These are the propositions that have moral content but no descriptive content. Finally, let us define an *impure moral proposition* as a proposition that has moral content but that is not a pure moral proposition.<sup>1</sup>

[So, for example, 'You acted wrongly in stealing that money' is an impure moral proposition, since it entails 'You stole that money.' But 'Stealing money is wrong' is a pure moral proposition, because it entails no such thing \(cf. Scanlon 2014, Ch. 2\).](#)

We are now in a position to see how the truth of an impure moral proposition could explain our belief that it is true. Consider, for example, the proposition that Pol Pot acted wrongly. This is clearly an impure moral proposition, since it has some descriptive content (it implies, for example, that Pol Pot did something or other). Suppose we begin with the belief that it's always wrong to order genocide. Then we encounter film footage of Pol Pot ordering genocide. Hence we come to believe that Pol Pot ordered genocide and that doing so is wrong. And from this we infer that Pol Pot did something wrong. In this case, while it may be true that Pol Pot's doing something wrong resulted in our believing that he did something wrong, it did so through the mediation of a descriptive belief, namely, the belief that he ordered genocide. From this belief, together with our purely moral belief that genocide is always wrong, we are able to infer the impure moral proposition that Pol Pot acted wrongly.

Note, however, that we cannot acquire our belief in pure moral propositions in this manner.

---

<sup>1</sup> To give a more precise account of impure moral propositions, we could define a *complete descriptive theory* as a maximally specific descriptive characterization of the world, and we could define a *total theory* as a conjunction of a complete descriptive theory and a complete moral theory. Then, just as we define a *pure moral proposition* as a disjunction of complete moral theories, we may define a *pure descriptive proposition* as a disjunction of complete descriptive theories. Finally, we may define an *impure moral proposition* as a disjunction of total theories that is not equivalent either to a pure moral proposition or to a pure descriptive proposition.

While a purely moral proposition that we antecedently believe together with a descriptive proposition that we come to believe may together entail an *impure* moral proposition (such as the proposition that Pol Pot acted wrongly), they cannot together entail a stronger pure moral proposition. Hence, if the clearest cases in which we come to believe a moral proposition *p* because *p* is true are cases in which we infer *p* from some descriptive proposition, together with our background moral beliefs, then we can get around the problem we have considered by restricting our attention to pure moral propositions.

Thus, we can avoid the problem by replacing the problematic premise (2\*) with the following:

(2\*\*) For any *pure* moral proposition *p*, it is not the case that our belief that *p* is explained by the fact that *p*, or vice versa, nor is it the case that the fact that *p* and our belief that *p* share a common explanation.

This premise, together with the Explanatory Connection Condition, entails that we can have no knowledge of *pure* moral propositions. But the skeptic wants more than this. She wants to maintain that we can have no knowledge of any moral propositions. To get this result, we will need one further assumption.

### 2.3 The Bridge Principle

The further assumption we need is the following:

**Bridge Principle:** For any person *S* and moral proposition *M*, if *S* knows that *M*, then there is some pure moral proposition *P* such that *S* is in a position to know that *P*.

The motivation for this principle is fairly straightforward. Assume the antecedent of the conditional is true. That is, assume that, for some arbitrary person *S* and moral proposition *M*, *S* knows that *M*.

Either *M* is a pure moral proposition or *M* an impure moral proposition. Suppose it's a pure moral proposition. In this case, there is clearly a pure moral proposition that *S* is in a position to know, namely *M* itself, since, obviously, if *S* knows that *M*, then *S* is in a position to know that *M*. Hence, if *M* is a pure moral proposition, then the consequent of the conditional is true.

Suppose, on the other hand, that *M* is an impure moral proposition, that is, a proposition whose

content is partially moral and partially descriptive. Suppose, for concreteness, that M is the proposition that Pol Pot acted wrongly. How might S know that M is true? It seems there are three possibilities: she might know it on empirical grounds, by learning certain descriptive facts about the world; she might know it on a priori grounds; or she might know it on the basis of a combination of a priori and empirical grounds. Regardless of which of these possibilities obtains, it must be that the totality of S's descriptive information about the world, together with whatever a priori grounds she may possess, constitutes a sufficient basis for knowing that Pol Pot acted wrongly. Let C be the conjunction of all the descriptive information S possesses. Thus, C will include all the information C possesses about the descriptive facts relevant to the rightness or wrongness of Pol Pot's actions, such as how many people he killed, the circumstances under which he killed them, etc. Now if, given the relevant a priori grounds, learning that C is true in the *actual world* suffices for S to know that, in the actual world, Pol Pot acted wrongly, then, given these same grounds, learning that C is true in any arbitrary world w should suffice for S to know that, in world w, Pol Pot acted wrongly. Hence, S must be in a position to know the following proposition:

P: In any world in which C is true, Pol Pot acts wrongly.

Note, however, that P is a pure moral proposition: it is equivalent to the disjunction of all the complete moral theories that entail that Pol Pot acts wrongly in worlds in which C is true. Hence, if S knows that Pol Pot acted wrongly, then S is in a position to know a pure moral proposition.

Recall that the proposition under consideration (that Pol Pot acted wrongly) was chosen simply for illustration. The same argument could be run for any impure moral proposition. And so we may conclude that if S knows any impure moral proposition, then S is in a position to know a pure moral proposition. And from this conclusion, together with what we have already established, the Bridge Principle follows.

## **2.4. Putting It All Together**

Combining these elements, we arrive at the following argument:

## Explanatory Trilemma Argument

1. (Connection Condition) For any proposition P, S knows that P only if there is some kind of explanatory connection between the fact that P and S's belief that P. In particular, if S knows that P, then either (A) the fact that P explains S's belief that P, (B) the fact that P is explained by S's belief that P, or (C) the fact that P and S's belief that P share a common explanation.
2. Therefore, for any proposition P, S is in a position to know that P only if S is in a position to form a belief that P such that either (A) this belief is explained by the fact that P, (B) this belief explains the fact that P, or (C) this belief and the fact that P share a common explanation.
3. For any pure moral proposition P, it is not the case that we are in a position to form a belief that P such that this belief is explained by the fact that P.
4. For any pure moral proposition P, it is not the case that we are in a position to form a belief that P such that this belief explains the fact that P.
5. For any pure moral proposition P, it is not the case that we are in a position to form a belief that P such that this belief and the fact that P share a common explanation.
6. Therefore, for any pure moral proposition P, it is not the case that we are in a position to know that P.
7. (Bridge Principle) For any moral proposition M, if S knows that M, then there is some pure moral proposition P such that S is in a position to know that P.
8. Therefore, for any moral proposition M, we do not know that M.

The overall strategy of this argument is to divide and conquer. Knowledge of any fact requires explanatory connections between belief and fact (Premise 1). And knowledge of *pure* moral propositions is ruled out because this requirement cannot be satisfied for pure moral propositions (Premises 3-5). Then, because pure and impure moral knowledge must stand or fall together (Premise 7), the unavailability of pure moral knowledge entails the unavailability of any moral knowledge.

### 3. Objections and Replies: A sketch

A number of objections may be raised to the argument just presented, and we will conclude by examining some of these objections and sketching the beginnings of some replies. The objections and responses listed here are not supposed to be an exhaustive categorization of the ways one can object to

or defend the Explanatory Trilemma Argument, any more than the survey in Section 1 was supposed to be an exhaustive survey of all possible arguments for moral skepticism. Nor are any of these arguments intended to be conclusive. This discussion should make clear what the main points are for and against the Explanatory Trilemma Argument, and thus help frame the debate over whether the Connection Concern can be developed into a sound argument for moral skepticism.

One response to the Explanatory Trilemma Argument rejects premise 1 by drawing a distinction between *a priori* and *a posteriori* knowledge. When it comes to empirical propositions that can only be known *a posteriori*, our knowledge of them requires sense-perception. Hence it involves forming beliefs in a manner that is influenced by, and hence explainable in terms of, the facts. But when it comes to necessary truths that are knowable *a priori*, such as mathematical truths, no such explanatory connections are required. And pure moral truths are like mathematical truths in that they are necessary and knowable *a priori*. Hence, pure moral knowledge likewise requires no such explanatory connections.

But to this objection there is a familiar reply. While mathematical truths may not be the right sorts of things to play a *causal* role, they may nonetheless play an explanatory role (Putnam 1972; Baker 2005; McGrath 2014). Thus, mathematical truths about derivatives and integrals may figure in an explanation of projectile motion. Indeed, it is not implausible to suppose that the fundamental principles of mathematics figure at least implicitly in many, if not all, scientific explanations. Moreover, it seems that all mathematical truths, or at least all the knowable ones, will be explainable in terms of the fundamental mathematical truths. Since such fundamental mathematical truths will plausibly figure both in the explanation of the *truth* of any given mathematical proposition and in the explanation of our *belief* in this proposition, it is plausible to suppose that our mathematical beliefs are explanatorily connected to the mathematical facts. Consequently, there is reason to be suspicious of the partners in guilt response.

But the argument just given cuts both ways. If we maintain that mathematical truths, though

necessary, can play a role in the explanation of our mathematical beliefs, why can't the moral anti-skeptic say the same thing about pure moral truths? Why can't she maintain that pure moral truths can figure in the explanation of our pure moral beliefs? Indeed, some philosophers have argued precisely this. One way to defend this claim is to maintain that moral truths are *analytic* or *conceptual* truths. It has been argued that, for any proposition P, if P is an analytic truth, then P might figure in the explanation of why anyone who possesses the concepts required to entertain P would be disposed to believe that P. And this fact, in turn, could figure in the explanation of our belief that P (Setiya 2014, p. 100ff).

It has been suggested that something similar may be true of purely moral propositions (Huemer 2005, p. 125-127; Cuneo and Shafer-Landau 2014). If "torture is wrong" is an analytic or conceptual truth, then perhaps it could figure in the explanation of our belief that torture is wrong. This would allow the anti-skeptic to reject premise 3 of the Explanatory Trilemma Argument, and hence resist the skeptical conclusion. This response, however, will only work if pure moral propositions are analytic or conceptual truths. And while this view has some defenders, it has far more detractors. To most philosophers, it seems that a person could have mistaken moral beliefs about the permissibility of abortion or euthanasia or the like while fully understanding the concepts that figure in these beliefs (Moore 1903).

Given the way we have defined the pure moral propositions, any pure moral proposition will be necessary. But it doesn't follow that they are *analytic*. Perhaps the pure moral facts are *facts about essences*, akin to the fact that all water is composed of hydrogen and oxygen, or the fact that all human beings are mammals. Such facts appear to be necessary without being analytic—we learn them by scientific inquiry rather than by reflecting on our folk concepts. Moreover, such facts about essences appear to be facts that can figure in scientific explanations, and are therefore the right kind of fact to explain our beliefs. Thus, if pure moral truths could be understood as facts about essences, then perhaps we could understand how they could be necessary, synthetic, and explanatory of our moral



beliefs. This would be another way of rejecting the third premise of the Explanatory Trilemma Argument.

As an illustration of a view of this kind, one might identify the fact that, e.g., *torture is wrong*, with the fact that a certain kind of attitude (say, moral disapproval) is fitting toward torture. And one might in turn identify the fact that *moral disapproval is fitting toward torture* with the fact that *torture possesses that natural feature which it is the functional role of moral disapproval to track* (Ross 2015). In this way, one would identify the pure moral fact that torture is wrong with the fact that torture stands in a certain relation to the functional role of moral disapproval. And facts about the functional role of mental state types are plausibly necessary truths about the essential nature of these types. Moreover, it seems these facts could play a role in explaining our moral beliefs. Suppose it were the functional role of the attitude of moral disapproval to track a certain feature, F, and that torture, by its very nature, possesses F. This could explain why we are disposed to have an attitude of moral disapproval toward torture. And this fact, in turn, could explain why we believe that torture is morally wrong. (This is not the only way such an account might go – see Copp (2008) for another account which holds that moral facts are identical with certain natural facts that plausibly explain our moral beliefs.)

Such an approach would, however, need to tackle the problem of moral disagreement. If there really is some natural feature which it is the functional role of moral disapproval to track, then one would expect there to be a high degree of uniformity in our dispositions to morally disapprove, and so it may be difficult, on this view, to explain the seeming ubiquity of moral disagreement.

So far we have focused on attempts to reject 3, the premise that denies that we are in a position to form moral beliefs that are explained by the pure moral facts. But some philosophers have attempted to reject 4 or 5, the premises that deny other kinds of explanatory connection between our beliefs and the pure moral facts. First, consider the view that the explanatory connection runs from our beliefs to the pure moral facts. On this view, we can know the moral facts, not because our beliefs depend on these facts, but because these facts depend on our beliefs. A view of this kind is defended by Sharon Street

(2006), who proposes that "the truth of 'X is a reason for agent A to Y' is a function of whether that judgment would be among A's evaluative judgments in reflective equilibrium." On this view, whether X is a reason for me to act in some way depends on whether I would regard it as such if my evaluative judgments were in reflective equilibrium, which in turn depends on my actual evaluative judgments. Hence, facts about my reasons depend, indirectly, on my beliefs about such facts.

But there are two reasons to doubt that such an approach can provide an adequate response to the moral skeptic. First, even if it could help to ground moral knowledge, it would be limited to what we might call egocentric moral knowledge, or knowledge of one's own moral reasons. For even if we adopt an internalist view on which my reasons depend on my evaluative beliefs, other people's reasons still won't depend on my evaluative beliefs--it simply isn't plausible, for example, that whether Julius Caesar had reason to cross the Rubicon depends on my evaluative beliefs. And so even on this view there won't be mind-to-world explanatory connection between my beliefs and other people's reasons that would allow me to have knowledge of the latter. Moreover, since moral reasons are generally thought to be universal, a view on which I can know my own moral reasons, without being in a position to know other people's moral reasons, may not be coherent.

The second problem with this approach is that it doesn't seem to apply to knowledge of pure moral facts. For pure moral facts are necessary truths. (This follows from the widely accepted claim that moral supervenes on the descriptive. For this claim implies that the true complete moral theory is necessarily true. And from this it follows that any true pure moral proposition (or disjunction of complete moral theories) must be necessarily true.) But since our evaluative judgments are contingent, any fact that depends on these judgments must likewise be contingent. And so the pure moral facts can't depend on our evaluative judgments. And while some moral anti-realists might defend an account of pure moral claims where even these conditional principles are contingent and dependent on individual attitudes, this must be counted as a substantial theoretical cost for the view.

Lastly, let us consider the rejection of 5, the premise that denies that our moral beliefs and the

pure moral facts can share a common explanation. A number of philosophers have proposed “Third Factor Accounts” according to which there is such a common explanation. This idea was first proposed in Nozick 1981, and has recently been advocated in Berker 2014, Enoch 2011, and Wielenberg 2010, *inter alia*. The basic idea is that there is some factor X which both explains why certain actions are morally required and explains why we believe they are morally required. Since many of these authors are responding to evolutionary debunking arguments, and accept, at least for the sake of argument, that our moral beliefs are shaped by evolutionary pressures, this third factor is typically identified with an evolutionary aim, such as “inclusive fitness” (Nozick) or “survival or reproductive success” (Enoch).

But there are several problems with this approach. One problem is that it clearly can only apply to *derivative* moral truths, or moral truths that have explanations. Many philosophers (including, we suspect, many philosophers who endorse Third Factor accounts), believe that there are some fundamental moral truths that explain other moral truths but that are not susceptible to further explanation. “Pain and suffering are bad” might be one example. “Survival and reproductive success are good” might be another. But if these, or any other, moral propositions cannot be explained, then clearly there can be no third factor that explains both them and our belief in them.

Moreover, even in the case of non-fundamental moral facts, it is dubious that the very same factor that explains our belief in them could also explain their truth. To illustrate the difficulty, suppose you believe that promise-keeping is morally required. What might explain why you believe this? Perhaps you believe this, at least in part, because dispositions to form this belief promote reproductive success. But *whose* reproductive success? From an evolutionary point of view, it’s irrelevant whether these dispositions promote *your* reproductive success. Rather, what matters is that they promoted the reproductive success of your ancestors—chiefly, your paleolithic ancestors. But it’s questionable whether this could have any bearing on whether promise-keeping is now morally required.

Moreover, even if we grant, with proponents of this approach such as David Enoch, that the evolutionary aim that explains the formation of our beliefs is good, this would at most explain why our

moral beliefs, or the dispositions that gave rise to them, are *good*—it would not explain why they are *true*. After all, false beliefs, like true beliefs, could promote good aims. This is an issue that utilitarians, beginning with Sidgwick (1907), have been acutely aware of. As Sidgwick pointed out, even if utilitarianism is true, and hence what is good is total happiness, the moral beliefs that most promote this good may be non-utilitarian, and hence false, beliefs. And the same may be true if what is good is survival or reproductive success. Suppose the aim I ought to promote is reproductive success, and hence that the valuable actions and valuable moral beliefs are the ones that promote reproductive success. In this case, it might turn out that I ought to cheat when I can get away with it (since doing so is conducive to reproductive success). But it might also turn out that I am disposed to believe that I ought *not* to cheat when I can get away with it, since having this false belief may be conducive to reproductive success (this could be true, for example, because other people are able to discern, to some degree, what my moral beliefs are, and because they are more likely to cooperate with me if I have this false belief). Hence, even if we grant that the evolutionary aim is good, it is plausible that evolutionary pressures could result in the formation of systematically false beliefs, and hence that they could constitute an off-track influence.

Thus, while there are a number of philosophical views that seem to offer an escape from the Explanatory Trilemma argument by allowing us to reject one or other of its premises, each of these views faces difficulties.

#### **4. Conclusion**

Skeptical arguments can play either of two roles: they can lead us to question *whether* we have a given kind of knowledge, or they can lead us to question *how* we can have such knowledge. This is true, in particular, of the skeptical arguments we have considered in this chapter. Some people, philosophers or non-philosophers alike, may wonder whether we have any moral knowledge, and may see some of the arguments we have considered as providing grounds for a negative answer. There are others, however,

who are more Moorean in their outlook, and who will regard it as obvious that we possess moral knowledge, and hence as obvious that these arguments must go wrong somewhere. Even for them, however, these arguments should prove illuminating, since replying to them requires giving a defensible account of how our moral beliefs are connected to the moral facts. And this, as we have seen, is no easy task.

## References

- Baker, A. (2005). Are there genuine mathematical explanations of physical phenomena? *Mind*, 223-238.
- Berker, Selim (2014). Does Evolutionary Psychology Show That Normativity Is Mind-Dependent? In Justin D'Arms & Daniel Jacobson (eds.), *Moral Psychology and Human Agency: Philosophical Essays on the Science of Ethics*. Oxford University Press. 215-252.
- Clarke-Doane, Justin (2014). Moral Epistemology: The Mathematics Analogy. *Noûs* 48 (2):238-255.
- Chisholm, R. M. (1966). *Theory of Knowledge*, Englewood Cliffs, NJ: Prentice Hall.
- Copp, David (2008). Darwinian skepticism about moral realism. *Philosophical Issues* 18 (1):186-206.
- Cuneo, Terence & Shafer-Landau, Russ (2014). The moral fixed points: new directions for moral nonnaturalism. *Philosophical Studies* 171 (3):399-443.
- Enoch, David (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press.
- Gettier, Edmund (1963). Is Justified True Belief Knowledge? *Analysis* 23 (6):121-123.
- Goldman, Alvin (1967). "A causal theory of knowing." *The Journal of Philosophy* 64 (12), pp. 357-372.
- Harman, Gilbert (1977). *The Nature of Morality: An Introduction to Ethics*. Oxford University Press.
- Huemer, M. (2005). *Ethical Intuitionism*. Palgrave Macmillan.

Joyce, Richard (2006). *The Evolution of Morality*. MIT Press.

McGrath, Sarah. "Relax! Don't do it! Why moral realism won't come cheap." *Oxford studies in metaethics* 9 (2014): 186-214.

Moore, G. E. (1903). *Principia Ethica*. Dover Publications.

Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.

Parfit, Derek. *On what matters: volume one*. Vol. 1. Oxford University Press, 2013.

Pollock, John (1987). Defeasible Reasoning. *Cognitive Science* 11 (4):481-518.

Putnam, Hilary (1972). *Philosophy of Logic*. London, Allen and Unwin.

Ross, Jacob. (2015) From Moral Blame to Moral Wrongness. Presented at the Arizona Workshop in Normative Ethics.

[Scanlon, TM \(2014\). \*Being Realistic About Reasons\*. Oxford University Press.](#)

Shafer-Landau, Russ (2004). *Whatever Happened to Good and Evil?*. Oxford University Press.

Sidgwick, Henry (1907). *The Methods of Ethics*, London: Macmillan;

Sinnott-Armstrong, Walter (2006). *Moral Skepticisms*. Oxford University Press.

Sosa, Ernest (1999). How to defeat opposition to Moore. *Philosophical Perspectives* 13 (s13):137-49.

Street, Sharon (2006). A Darwinian dilemma for realist theories of value. *Philosophical Studies* 127 (1):109-166.

Sturgeon, Nicholas (1988). Moral Explanations. In *Essays on Moral Realism*, ed. Geoffrey Sayre-McCord. Cornell University Press.

Wielenberg, Erik J. (2010). On the evolutionary debunking of morality. *Ethics* 120 (3):441-464.

## **Related Topics**

Error theory, Reliability and explanatory efficacy challenges, Intuitionism in moral epistemology, The epistemic significance of ethical disagreement