BRILL

JOURNAL OF CHINESE PHILOSOPHY 49 (2022) 405–415

Journal of Chinese Philosophy

brill.com/jcph

Festschriften

Unexceptional Moral Knowledge

Timothy Williamson | ORCID: 0000-0002-4659-8672
Faculty of Philosophy, University of Oxford, Oxford, UK
timothy.williamson@philosophy.ox.ac.uk

Abstract

The article defends moral realism against epistemological objections by arguing that if there are moral truths, some of them are known. The claim that moral properties are unknowable because causally inert is shown to be ineffective: none of the main current theories of knowledge requires a causal connection, and anyway moral properties have not been shown to be causally inert. It is explained why a posteriori moral knowledge need not derive from combining a priori moral knowledge with a posteriori non-moral knowledge. The possibility of moral knowledge by perception and by testimony is briefly defended. The role of recognitional capacities for instances of moral properties is emphasized.

Keywords

moral realism – moral epistemology – moral perception – moral testimony – a priori – a posteriori

Opponents of moral realism often argue that, even if there were moral truths, we could not know them. That conditional excludes the most natural basis for asserting that *there are* moral truths: that we know some.[1]

Epistemological objections to moral realism typically ignore contemporary developments in epistemology. Some invoke the alleged causal inertness of moral properties, when most epistemologists reject a causal theory of knowledge. In this position paper, I sketch a view on which, if there are moral truths, there is no special obstacle to knowing them. Indeed, we *do* know some moral truths, and so have the most natural basis for asserting that there are such truths. For instance, we know that slavery is wrong.

This paper is argued on the *supposition* that there are moral truths. Of course that supposition trivially entails that it is *a fortiori possible* – theoretically and practically, metaphysically and

1 In November 2013, Professor Chung-ying Cheng and I had a public dialogue in Oxford as representatives of Chinese philosophy and Western philosophy respectively. We were gratified to find much common ground between our positions. In particular, we both endorsed a robust realism about the natural world, which extended to a robust moral realism. Thus, when he invited me to contribute an essay to the special issue of the *Journal of Chinese Philosophy* to celebrate its fiftieth birthday in 2023, I found it appropriate to write on moral realism. I thank Farbod Akhlaghi and Daniel Kodsi for useful comments on a draft of this paper.

epistemically – that there are moral truths. But the supposition does *not* entail that it is also *indubitable* that there are moral truths. While no truths are impossible, many truths are dubitable. In supposing that there are moral truths, I do not suppose away metaethical disagreement.

Elsewhere, I have argued that salient objections to moral truth rest on metaphysical and semantic confusions.[2] I do not invoke those arguments here. For present purposes, the "if" in "if there are moral truths" remains.

My focus will be on *pre-theoretic* moral knowledge, available to ordinary human agents, not on the sorts of theoretical moral knowledge that moral philosophers may hope to achieve. We could hardly obtain the latter without first having some of the former, just as we could hardly obtain theoretical physical knowledge without first having some pre-theoretic physical knowledge, available to ordinary human agents through sense perception.

1

In assessing putative moral knowledge, one must take care to apply normal standards for knowing, not specially raised standards whose general application would imply widespread scepticism even outside the moral domain. For, in this dialectical context, widespread non-moral knowledge is granted: the issue is whether, if there were moral truths, they would be at some special disadvantage in epistemic accessibility, compared to ordinary non-moral truths. By non-sceptical standards, we know plenty about our environment, even though we have plenty of false beliefs about it too. In favourable circumstances, even fallible cognitive capacities can generate knowledge. Thus one cannot refute claims of moral knowledge simply by describing actual or possible social groups with contrary moral beliefs – false beliefs, according to the moral realist. Our fallibility on moral issues is quite compatible with having plenty of moral knowledge, by non-sceptical standards for knowledge.

For example, objectors to the claim to know that slavery is wrong say: "You could have been brought up in a social group where the belief that slavery is morally unobjectionable was deeply rooted, and anyone who questioned it was regarded as crazy. In those circumstances, you too would have believed that slavery is morally unobjectionable." The objector is using the envisaged society as a *sceptical scenario* for beliefs about the moral status of slavery. The background assumption is that ordinary beliefs on the matter are formed by a process of socialization, and the point of the scenario is to show that, if our belief is true, such a process could equally have led to a false belief on the same question. A further background assumption is that the actual moral status of slavery remains constant between our actual society and the imagined one, which most moral realists will grant (a sceptical possibility where people falsely believe that slavery is wrong is morally impossible). The objector concludes that we don't know that slavery is wrong, because our belief is formed by an unreliable process.

By contrast, all parties agree that we know that the earth is round. Someone might object by saying: "You could have been brought up in a society where the belief that the earth is flat was deeply rooted, and anyone who questioned it was regarded as crazy. In those circumstances, you too would have believed that the earth is flat." The objector is using the envisaged society as a sceptical scenario for beliefs about the earth's shape. Again, the background assumption is that ordinary beliefs on the matter are formed by a process of socialization, and the point of the scenario is to show that, if our belief is true, such a process could equally have led to a false belief on the same question. A further

2 Timothy Williamson, "Moral Anti-Exceptionalism", in *The Oxford Handbook of Moral Realism*, eds. Paul Bloomfield and David Copp (Oxford: Oxford University Press, forthcoming).

background assumption is that the real shape of the earth remains constant between our actual society and the imagined one, which all parties grant (a sceptical possibility where people falsely believe that the earth is round is scientifically impossible). The objector concludes that we don't know that the earth is round, because our belief is formed by an unreliable process.

By normal epistemic standards, the latter objection fails. We *do* know that the earth is round; the sceptical scenario is too "distant" from our actual situation to undermine our actual knowledge. But if the latter sceptical scenario does not achieve its intended purpose, why should we regard the former sceptical scenario as any more successful?

Opponents of moral realism will try to identify an epistemic asymmetry between the two cases. They may emphasize our reliance on scientific expertise in our knowledge that the earth is round. Of course, in the sceptical scenario, our belief that the earth is flat may also result from deference to people who play the role of scientific experts in that society; they may genuinely know much more than others. They dismiss any suggestion of a round earth as ridiculous and discredited, giving abstruse explanations of its impossibility.

Objectors respond that we rely on *genuine* expertise, whereas flat-earthers rely on mere pseudo-expertise, so the possibility of the flat-earther society does not exclude our knowledge that the earth is round, whereas the possibility of the slavery-tolerant society *does* exclude our putative knowledge that slavery is wrong. Thus they argue *from* the belief's epistemic status *to* the error possibility's (ir)relevance. That is not inappropriate in itself, but it raises an obvious question: why can't the moral realist argue *from* the status of our belief that slavery is wrong as knowledge *to* the irrelevance of the possible slavery-tolerant society? Just as the flat-earther society is less scientifically advanced than ours, so the slavery-tolerant society is less morally progressive than ours. Thus critics of moral realism may already presuppose that we don't know that slavery is wrong when they take the possibility of the slavery-tolerant society to exclude our putative knowledge that slavery is wrong.

In brief, critics of moral realism cannot rely on applying generic forms of sceptical argument to the moral case. They need to target their arguments more accurately. As already noted, they also cannot rely on the supposed causal impotence of the moral, since a causal theory of knowledge is not in play. At best, they could argue that, in the moral domain, causal impotence excludes some other property that is necessary for being known.

2

The metaphysical possibility of error does not justify scepticism. That holds generally, and, more specifically, in the moral domain. Links from error-possibilities to ignorance are *local*: only error-possibilities in close cases exclude knowledge. But what makes a case 'close'?

In a space of possible worlds or cases, comparative closeness represents comparative similarity of some relevant kind. For example, in a case similar to one in which S comes or continues to believe p in a way w, an agent S* similar to S comes or continues to believe a proposition p^* similar to p in a way w^* similar to w. Often, S* just is S, p^* just is p, and w^* just is w, but such identities are not generally required. The idea is then that a case of knowledge is not similar to a case of false belief. In other words, knowledge needs a *margin for error*, or amounts to *safety* (from error). The required level of similarity need not be constant across cases. This is not intended as a reductive definition of knowledge in terms of similarity, for we may be unable to specify the required levels of safety without invoking knowledge itself, implicitly or explicitly. Of course, proponents of a knowledge-first approach in epistemology do not expect such a reductive definition of knowledge anyway. Nevertheless, the picture of knowledge as safety from error can still

be of use, by highlighting key structural features of knowledge.[3]

For present purposes, the point is that, once we grant that moral beliefs can be *true* (as we do in this dialectical context), we have no special reason to doubt that they can also be *safe*. A case of true moral belief is a point in logical space; a case of safe moral belief requires only a small region of such points. For example, a case where we truly believe that it is wrong to torture a political prisoner is not typically close to anything like a case where we *falsely* believe that it is wrong to torture a political prisoner. Thus, on a safety conception of knowledge, bridging the gap between the possibility of moral true belief and the possibility of moral knowledge looks quite feasible. To deny moral knowledge while for the sake of argument granting moral truth is an ill-motivated position.

In accepting moral knowledge, the safety theorist is not committed to the causal efficacy of moral properties. Still, on a modal conception of causality, perhaps in terms of counterfactual conditionals, resisting the causal efficacy of moral properties seems ill-motivated. If our moral beliefs correlate well with moral truth in some region of logical space, why shouldn't the wrongness of an action have the counterfactual correlations with our belief in its wrongness required for the wrongness to cause the belief?

Opponents of moral realism may argue that what *really* caused the moral belief was not the wrongness itself but the underlying physical facts. But such an argument betrays a naïve understanding of causality. First, even if the underlying physical facts did cause the belief, it does not follow that the wrongness did not also cause the belief; most effects have many causes, at different levels of generality. Second, the wrongness may even turn out to be a *better* candidate cause for the belief than the underlying physical facts, because the former is more strongly correlated than the latter with the putative effect. When the underlying physical facts are presented as a rival cause, a fairly specific physical description is imagined, to achieve causal sufficiency. But such a specific putative cause will fall far short of the putative effect in generality: even had the physical facts been slightly different, there would still have been the same moral belief. By contrast, the moral fact (the wrongness) may easily have a generality far more like that of the putative effect (the belief in wrongness). Amongst the available candidates, the moral fact may well have the strongest modal correlation with the moral belief.[4] Of course, how causality relates to modal correlation is itself a controversial issue, but what is currently understood about causality provides no good basis for denying that an act's being wrong can cause someone to believe that it is wrong.

Irrespective of the causal issue, a safety model of knowledge provides a hospitable setting for moral knowledge, compatible with many different implementations in detail of the approach.

Other forms of epistemological locality are also unthreatening to moral knowledge. For example, on a standard form of *contextualism* about the word 'know', it correctly applies in a given context of utterance when all *contextually relevant* error-possibilities are excluded. In an ordinary conversation about the moral status of something someone has just done, radical possibilities of error in participants' shared moral convictions are not treated as relevant. Consequently, unless those standards are in fact erroneous, the contextually relevant standard for knowledge may well be met. Thus contextualism too can facilitate bridging the

3 For margins for error see Timothy Williamson, "Inexact Knowledge", *Mind*, 101 (1992): 217–242. For safety conceptions of knowledge see Ernest Sosa, "Skepticism and Contextualism", *Philosophical Issues*, 10 (2000): 1–18 and Timothy Williamson, *Knowledge and its Limits* (Oxford: Oxford University Press, 2000). Williamson (ibid.) applies a knowledge-first approach to epistemology. For further development and explanation of these ideas see Timothy Williamson, "Replies to Critics", in Patrick Greenough and Duncan Pritchard eds., *Williamson on Knowledge* (Oxford: Oxford University Press, 2009) and "Probability and Danger", *The Amherst Lecture in Philosophy*, 4 (2009): 1–35.

4 Williamson, *Knowledge and its Limits*, 80–88 develops a similar approach in terms of correlation coefficients.

gap between the possibility of moral true belief and the possibility of moral knowledge.[5]

The focus on objective error-possibilities will not satisfy some epistemological *internalists*. In their view, no level of objective reliability suffices to justify a belief, because that dimension is not itself accessible from the perspective of the subject's consciousness. The canonical example is a person who in fact has a 100% reliable power of clairvoyance, but without internal corroboration.[6] Dialectically, the example is ineffective, because externalists can respond that its presentation confuses the relevant case, where clairvoyance functions like a normal unreflective sensory power, with the distracting case – to which the context of epistemological debate calls attention – where the subject reflects on the strange presentation and comes to believe its content on the quite different basis of a shaky abductive inference.[7] Fortunately, the issue need not be settled here. For internalist constraints pose no special obstacle to moral knowledge or the justification of moral belief. Purely from the perspective of the subject's consciousness, a full-hearted moral conviction, coherently combined with other beliefs and seemings, may be as good as belief gets. Any distinctive threat to moral knowledge or the justification of moral belief concerns the dimension of objective reliability. Thus, for present purposes, we can focus on an externalist conception of knowledge.

Rather than try the reader's patience by attempting a comprehensive survey of contemporary theories of knowledge, I will make a generalization: they offer no distinctive, credible threat to moral knowledge. There are various general sceptical threats to knowledge, but they are neither distinctive to moral knowledge nor seriously credible.

3

Another form of attack on moral knowledge consists in challenging moral realists to explain *how* they know moral truths. Their opponents suspect that sooner or later moral realists will have to fall back on a suspiciously convenient faculty of *moral intuition*.

An example: I judge "Poisoning Alex was wrong". Is that a moral intuition? I know that Alex was poisoned, from various reliable news sources; it is a normal case of knowledge by testimony. I then make the moral judgment that poisoning him was wrong. Although I did not reach my conclusion by conscious step-by-step reasoning, I can provide non-deductive support for it, by citing known circumstances of the case and other considerations. In those respects, it is not so different from my non-moral judgment "Poisoning Alex was premeditated". In that case too, although I do not reach my conclusion by conscious step-by-step reasoning, I can provide non-deductive support for it, by citing known circumstances of the case and other considerations. In each case, I applied a new term ("wrong", "premeditated") in the light of my evidence. Presumably, I *can* thereby come to know that poisoning him was premeditated. If my judgment "Poisoning him was premeditated" thereby depends on "intuition", presumably so too does my judgment "Poisoning him was wrong", but since one can achieve knowledge by depending on intuition, as in the other case, there is no obvious cause for alarm.

More generally, we have *recognitional capacities* for properties and kinds of many types, enabling

5 For contextualist approaches see, for example: Gail Stine, "Skepticism, Relevant Alternatives and Deductive Closure", *Philosophical Studies*, 29 (1976): 249–261, Stewart Cohen, "How to be a Fallibilist", *Philosophical Perspectives*, 2 (1988): 91–123, Keith DeRose, "Contextualism and Knowledge Attributions", *Philosophy and Phenomenological Research*, 52 (1992): 913–929, and *The Case for Contextualism: Knowledge, Skepticism, and Context, Volume 1* (Oxford: Clarendon Press, 2009), and David Lewis, "Elusive knowledge", *Australasian Journal of Philosophy*, 74 (1996): 549–567.

6 Laurence BonJour, *The Structure of Empirical Knowledge* (Cambridge, Mass.: Harvard University Press, 1985).

7 Jennifer Nagel, "Knowledge and reliability", in *Alvin Goldman and his Critics*, eds. Hilary Kornblith and Brian McLaughlin (Oxford: Blackwell, 2016).

us to recognize whether a given case instantiates them. Most people have recognitional capacities for many species of plants and animals and for many types of artefact – types of clothing, types of furniture, types of household utensil, types of vehicle, types of building, and so on. They can recognize various types of weather, various types of art and music, and various types of behaviour: hasty or leisurely, careless or careful, confident or timid, rude or polite, cold or warm, serious or humorous, hostile or friendly, cruel or kind. When we apply a recognitional capacity, we do not typically use conscious step-by-step reasoning, but our judgment is still evidence-based. To describe recognitional capacities in all those cases as based on a faculty of 'intuition' merely obfuscates what sort of pattern recognition is going on.

A different view of particular moral judgments like "Poisoning Alex was wrong" is widespread in moral epistemology. On that view, such judgments are *inferential*. They come from arguments like this, where "D" is schematic for a qualitative *non-moral* description:

Premise 1	Poisoning Alex was D.
Premise 2	Whatever is D is wrong.
Conclusion	Poisoning Alex was wrong.

The argument is deductively valid, with Premise 2 read as a universal generalization over all times. Premise 1 is particular, purely non-moral, and *a posteriori*; Premise 2 is general, moral, and supposedly *a priori*; the Conclusion is particular, moral, and *a posteriori*. On this view, the moral element comes from a capacity to assess general moral principles *a priori*. By contrast, on the view suggested before, the moral element comes from a recognitional capacity to assess particular cases morally *a posteriori*.[8]

The inferential view of moral judgment is implausible, both psychologically and epistemologically. Psychologically, we are not aware of making any such deduction, or even of entertaining any such universal generalization as Premise 2; nor have we any evidence that what goes on in us unconsciously takes any such artificial form. Epistemologically, we are typically *more* confident of the Conclusion than of Premise 2, because doubts as to whether "D" excludes all potential counterexamples absent from the Alex case pose a threat to Premise 2 but not to the Conclusion. Thus, even if some of our confidence in the Conclusion comes from Premise 2, not all of it does. Formulating universal moral principles like Premise 2 takes us into a realm of proto-philosophical speculation that we did not enter merely by making a moral judgment on a particular case.

The envisaged inferential alternative fails for a deeper reason too. For where does Premise 2 itself come from? If we are in a position to have ordinary pre-theoretic knowledge of Premise 1, the non-moral description 'D' will be couched in reasonably accessible, not too complicated terms. But such a description is likely to make Premise 2 exception-prone and so false, even in our own eyes – poisoning Hitler in 1939 would not have been wrong, and so on. Premise 2 is hardly innate, nor were we taught it as children.

Will moral realists invoke "intuition" here, to explain our putative knowledge of Premise 2? That is what their opponents suspect. What is most dodgy about such reliance on "intuition" is not that it sounds spooky but that it fails to engage with the semantic structure of the sentence. Premise 2 is a universal generalization of the form "Whatever is F is G". Normally, to recognize such a sentence as true, without matching it to a list of pre-given truths, one must recognize some necessary or contingent connection between the two properties, being F and being G (here, being D and being wrong). One can articulate that connection as linking the simpler sentences "*x* is F" and "*x* is G", with the variable "*x*" for an unspecified instance, prior to introducing the universal quantifier ("whatever"). More specifically, by the informal analogue in natural language of a standard introduction rule for the universal quantifier in a system of natural deduction, having verified "*x* is G" conditionally

8 For an account of such a form see Christopher Peacocke, *The Realm of Reason* (Oxford: Clarendon Press, 2004): 198–231.

on the supposition "*x* is F", one can then verify the universal generalization "Whatever is F is G". That is logically the most basic way of verifying a universal generalization of that form. To verify Premise 2, it means supposing an unspecified instance "*x* is D" and verifying "*x* is wrong" under that supposition. But "*x* is wrong" is just an unspecified instance of the same pattern as the Conclusion, 'Poisoning Alex was wrong'. Moreover, making a judgment under a supposition, imaginatively, is just the offline analogue of making the judgment unconditionally, "live" online, in a particular case. To make such a judgment we need just the sort of recognitional capacity (still for "wrong") invoked by the simpler account above.[9]

In brief: assenting to a sentence on the basis of intuition was not intended to bypass semantic understanding of that sentence, but instead to work through it; the most natural way of applying that approach to Premise 2 makes the inferential apparatus redundant.

Might our assent to Premise 2 come from some more roundabout process of reasoning? One could postulate that we derived Premise 2 deductively from more basic moral principles, but that only postpones the general problem to those latter principles, on pain of an infinite regress: how do we know the more basic moral principles? Similarly, to postulate that we arrived at Premise 2 inductively or abductively returns us to the original problem, since the inductive or abductive evidence for Premise 2 would presumably itself include further particular moral claims like the Conclusion.

We could try replacing the universal generalization "Whatever is D is wrong" as Premise 2 by a merely generic generalization or *ceteris paribus* "law", "D actions are wrong", making the argument non-deductive. But that modification does not address the underlying problems, which arise in similar ways for the new Premise 2. Without the recognitional capacity for particular cases, we still cannot verify the generic generalization.[10]

The upshot of these considerations is that inferential accounts fail to present a genuine alternative, since they will not work without a recognitional capacity, the need for which they merely obscure. Their extra complexity brings no extra explanatory power.

Of course, recognitional capacities themselves involve hidden complexity, like the neural processes underlying face recognition. Any recognitional capacity for a moral property, however limited, will involve such hidden complexity too. But the key objection to the inferential account is *not* that it postulates hidden complexity. Rather, it is that the inferential account is *regressive*. For it analyses the cognitive process by which we make a singular ascription of "wrong" (typically a novel one), our assent to which itself depends on a prior singular ascription of "wrong", albeit offline. If every singular ascription of "wrong" depended on a prior singular ascription of "wrong", we could never get started.

By contrast, applying a recognitional capacity involves no such regress. No component of the neural processes underlying face recognition is itself another recognitional process as complex as the one we started with (a homunculus recognizing a face on the screen of a mental cinema). The same applies to our recognitional capacities for moral properties.

Once we have a recognitional capacity for "wrong", we can use it to start formulating general moral principles and testing them against particular cases, real and imagined. We may adopt some of those principles, and even start using them to correct the deliverances of the original capacity, where we detect a bias or anomaly. We may indeed make some singular ascriptions of "wrong" inferentially, using those principles as premises, as with

9 For the suppositional assessment of universal generalizations see Timothy Williamson, *Suppose and Tell: The Semantics and Heuristics of Conditionals* (Oxford: Oxford University Press, 2020): 142–6. For the connection between offline and online judgments see Paul Boghossian and Timothy Williamson, *Debating the A Priori* (Oxford: Oxford University Press, 2020): 121–2 and 179–81.

10 The account of the suppositional assessment of generalizations in Williamson, *Suppose and Tell* (2020), is applied to generics.

the argument from Premise 1 and Premise 2 to the Conclusion. That would not be regressive. What triggers the regress is the assumption that *every* singular ascription of "wrong" is inferential.

That we have recognitional capacities for moral properties is no mere postulate of moral realists. Much ordinary moral discourse takes it for granted, for example when we describe someone as recognizing that an action was wrong. Phenomenologically, much ordinary moral judgment (and misjudgement) feels like the exercise of recognitional capacities for moral properties.

Naturally, a philosopher can still insist that what we have are *failed* recognitional capacities. For a moral anti-realist, there are no moral properties to be recognized. For an error theorist, the would-be recognitional capacity is like one for the property of being bewitched. But such views are irrelevant to the dialectic of this paper. By hypothesis, there are moral truths, for the question is whether a distinctive problem of moral knowledge still arises under that supposition. It is a short step from moral truths to moral properties within a standard semantic framework, since moral truths require moral predicates to have non-empty intensions, and those intensions are *ipso facto* moral properties in the relevant sense.[11] Thus we may assume that there are moral properties like being wrong for the associated recognitional capacities to track.

Of course, not all recognitional capacities are good at their job. Someone could hold that, although the word "wrong" does express the moral property of being wrong, the associated recognitional capacity is bad at tracking that property, indeed, is hopelessly unreliable. Some moral realists take our would-be recognitional capacities for moral properties to have been definitively superseded by the clear principles of their favoured moral theory – Kantianism, utilitarianism, whatever. However, such views pose little threat to the present strategy. They hold that one can do much better epistemologically by theoretical reflection on moral principles than by reliance on common sense recognitional capacities. Indeed, their scepticism about the latter is justified by the alleged epistemological superiority of the former when the two conflict. The suggestion is not that the moral domain is epistemically inaccessible, just that it is better accessed by theoretical reflection than by ordinary recognitional capacities. One may doubt that the theoretical reflections at issue are as cogent as claimed, but there is no need to press such doubts here, for either way the moral domain is epistemically accessible.

An alternative kind of argument for the unreliability of moral recognitional capacities invokes individual and cultural variation in their deliverances. Such reasoning grants the associated moral properties, at least for the sake of argument, and does not privilege some alternative form of epistemic access to the moral domain. The idea is just that if one person or group judges an action as wrong, while another judges it as right, one side is in error (since we are assuming moral realism, no relativistic doublethink is in play). Thus the extent of moral disagreement puts an upper limit on moral reliability.

Such arguments from disagreement are not illegitimate in principle, but must be applied with care. In particular, we must not forget the earlier lesson that moral knowledge requires only local reliability. For example, suppose that members of society S judge that actions of type A are wrong, while members of society S* judge that actions of type A are right. If members of each society restrict their judgment to actions performed within their own society, so far there is no disagreement, since whether an action of a given type is right or wrong may depend on the social setting: a physical movement may constitute an offensive gesture in S but not in S*. However, suppose that members of each society generalize to the other: members of S judge that instances of A are wrong, whatever

11 See Williamson, "Moral Anti-Exceptionalism" for more on moral realism within an intentional semantic framework, and Williamson "Epistemological Consequences of Frege Puzzles", *Philosophical Topics*, 49 (2021), 287–320. for the complications raised by the ascription of propositional attitudes – including knowledge – when propositions are individuated intentionally.

society they are performed in, while members of S* judge that instances of A are right, whatever society they are performed in. Now there is genuine disagreement, since the judgments are made about the same domain of actions. At least one side is in error somewhere. Nevertheless, it may still be that members of each society are reliable about whether A-actions performed in their own society are wrong: instances of A performed in S are wrong, while instances of A performed in S* are not wrong. We should not be surprised when people are better at understanding the moral significance of actions performed in their own society than of actions performed in another society, with which they are less familiar. Thus members of S may *know*, of particular A-actions performed in S, that they are wrong, while members of S* *know*, of particular A-actions performed in S*, that they are right.

That example is obviously very simple and schematic. Sometimes, the asymmetry may run the opposite way, with members of each society blind to its own faults but alert to the faults of its neighbours. The point is just that widespread moral disagreement is quite compatible with widespread moral knowledge of specific matters.

Gestures at the extent of moral disagreement are often quite perfunctory. In surveying the data, one danger to avoid is *cherry-picking*. For example, given any moral belief, one can probably find people somewhere with a contrary moral belief. But that shows nothing special about morality. Human individuals and human societies are very various. Given any *non-moral* belief, one can probably find people somewhere with a contrary non-moral belief. Disputes between proponents and opponents of the theory of evolution are no more likely to be resolved than disputes between proponents and opponents of a right to abortion. On both moral and non-moral matters, disagreement is typically louder and more attention-grabbing than agreement, making us liable to underestimate the latter's prevalence.

Still, there *is* plenty of moral disagreement. Some philosophers may therefore be tempted to dismiss local patches of correlation between moral belief and moral truth as products of chance rather than cases of moral knowledge. Wouldn't one expect such patches, if moral belief and moral truth varied independently of each other over a large enough domain? As just explained, that attitude may well underestimate the extent of moral agreement. But it also ignores another crucial factor: the *metasemantics* of moral terms. Like other words, "wrong" does not get its intension by magic. Speakers are willing to apply it in some cases and not in others; the patterns of their use help determine its reference. A principle of charity constrains correct interpretation. Arguably, the appropriate general principle maximizes the *knowledge* expressed by the relevant discourse: for "right" and "wrong" in their moral senses, that is moral knowledge.[12] Although the correct interpretation is the resultant of several forces, correlations between moral belief and moral truth are far from random: they manifest in part a constitutive pressure towards knowledge.

In short: one would expect that, if moral discourse can express moral truths, it does express much moral knowledge. General epistemological theory is quite consistent with that expectation.

4

We can briefly consider some specific kinds of moral knowledge that have proved controversial:

12 For the role of charity in interpretation see Richard Grandy, "Reference, Meaning, and Belief", *Journal of Philosophy*, 70 (1973): 439–52 and Donald Davidson, "The Method of Truth in Metaphysics", *Midwest Studies in Philosophy*, 2 (1977): 244–54. For a knowledge-maximizing principle of charity see Timothy Williamson, *The Philosophy of Philosophy* (Oxford: Blackwell, 2007): 248–78. For a reinterpretation of that principle in terms of an account of mindreading in humans and other animals, see Timothy Williamson, "Where did it come from? Where will it go?", in *Putting Knowledge to Work: New Directions for Knowledge-First Epistemology*, eds. Arturs Logins and Jacques-Henri Vollet (Oxford: Oxford University Press, forthcoming).

in particular, moral knowledge by *perception* and by *testimony*. They have sometimes been regarded as cases where moral epistemology diverges from non-moral epistemology. I will suggest that they are much less exceptional, and much less problematic, than they may sound.

Perceptual language is often used in an extended or metaphorical sense applicable to both sensory and non-sensory cognition, including moral cognition. Someone can *perceive* a research programme as flourishing, or come to *see* that it is degenerating. Likewise, someone can *perceive* slavery as acceptable, or come to *see* that it is wrong. Such examples are no more contentious than moral knowledge; leave them aside. The question is whether there can be moral knowledge by genuinely sensory perception. For example, can one gain moral knowledge *by seeing*, in a sense which requires one to use one's eyes, though of course also one's brain?

Seeing-that must be distinguished from object-seeing. One can see an Anglo-Saxon coin without seeing *that* it is an Anglo-Saxon coin, but an expert numismatist can see *that* it is an Anglo-Saxon coin. If we have an encapsulated visual module, impervious to background information, it will not by itself generate judgments like "That's an Anglo-Saxon coin". Nor, presumably, will it generate judgments like "That's a giraffe" or "That's Roger Federer". Perhaps it confines itself to shapes and colours. Indeed, by itself, it will generate very little of the seeing-that which we need in science and in ordinary life – potentially including the ordinary moral life.

In the moral case, the issue is not whether one can see special moral objects, but whether one can see *that* an object has a moral property. Sometimes, one can see *that* a kicking was deliberate. Sometimes, one can also see *that* it was wrong. Some philosophers may insist that one only sees that the kicking had various non-moral properties, and then infers that it was wrong by an argument like that in section 3 from Premise 1 and Premise 2 to the Conclusion. However, there is no good reason to impose such an inferential structure. Seeing-that is not restricted to contents internal to a primitive module for vision. Sometimes, an expert chess-player sees that a position is a win for black, and an expert boxer sees that his opponent is tiring. By other cues, sometimes a psychologically competent person sees that a kicking was deliberate, and a morally competent person sees that it was wrong.

Seeing-that is a way of knowing-that: when you see that something is so, you thereby know that it is so.[13] Thus, when you saw that the kicking was wrong, you thereby knew that it was wrong. Consequently, there is moral knowledge by perception.[14]

The case of moral knowledge by testimony may seem trickier, since it raises issues about the autonomy of the moral agent. Still, if a trusted colleague tells us that she had witnessed a job candidate behaving in a morally obnoxious way, but that confidentiality prevents her from explaining what it was, many of us would be willing to take her word for it.

In a purer case of moral testimony, the two parties already share all their relevant non-moral information, and the proffered testimony concerns solely the moral verdict on the case. For example, someone who trusts you comes to you with his moral problem, fully explains the relevant circumstances, and asks you what he ought to do. Perhaps you offer to talk through the options with him, while leaving him to make the final decision. He doesn't trust himself to make the decision by himself, he tells you, since he fears that self-interest will influence his judgment to an unknown degree. He would rather bind himself in advance to doing whatever you tell him he ought to do (your own self-interest is not at stake). He will simply accept,

13 Williamson, *Knowledge and its Limits*, 33–41.

14 See Sarah McGrath, "Moral Knowledge by Perception", *Philosophical Perspectives*, 18 (2004): 209–228 and *Moral Knowledge* (Oxford: Oxford University Press, 2019) for more on more moral knowledge by perception.

and act on, your testimony as to what he ought to do in this important matter (no pressure). You may not like the position in which he has put you, but does he have to be confused?

In that example, he simply treats you as better positioned to judge a particular moral issue, because you are not an interested party. You are not being treated as a general moral expert. Much scorn has been poured on the idea of moral expertise – perhaps partly from a suspicion that some moral philosophers want to set *themselves* up as moral experts, to whom non-experts should defer. But the hostility to moral deference may also manifest the individualism of some Western societies, by comparison with many non-Western societies: it is far from culturally universal. Imagine someone of wide human experience, who is often consulted in difficult conflicts, but otherwise does not interfere; her advice has generally led to outcomes that almost everyone regards in retrospect as fair and reasonable. It does not seem crazy to regard her as unusually *wise*, in a morally non-neutral sense. The thought 'She understands moral matters better than I do' is not intrinsically disreputable. Someone might think it, and decide to follow the wise woman's advice, without attempting to second-guess it in any way.

Some philosophers will find that deferential attitude abject, an abnegation of moral agents' responsibility for their own decisions. "Don't be so spineless!", they might say. But the deferential agent could reasonably respond: "It's not just about me; this decision impacts other people. My primary responsibility is to do right by them, not to indulge self-regarding concerns about my moral agency or autonomy. My best chance of doing the right thing is by following the wise woman's advice, no questions asked." In some circumstances, a willingness to defer to others on a moral question may be *morally* superior to an insistence on following one's own instincts. Of course, someone who defers like that is unlikely to understand *why* the chosen option is best. But, in the circumstances, *not* deferring is no route to such understanding either. Anyway, such lack of understanding is not specific to *moral* deference. Someone who merely defers on a purely prudential decision, no questions asked, is also unlikely to understand why the chosen option is best.

Knowledge by perception and knowledge by testimony are two of the most commonly cited cases where moral epistemology has been thought to differ from general epistemology. On closer inspection, the alleged differences vanish. This is further evidence that moral epistemology is much less special than it has been thought to be.[15]

5

On the evidence of this paper, if there is moral truth, there is also moral knowledge, of unexceptional kinds. Epistemological arguments against moral realism are based on bad epistemology.

15 For more on moral knowledge by testimony and moral expertise see: Sarah McGrath "Moral Disagreement and Moral Expertise", *Oxford Studies in Metaethics*, 3 (2008): 87–108, "The Puzzle of Pure Moral Deference", *Philosophical Perspectives*, 23 (2009): 321–344, "Skepticism about Moral Expertise as a Puzzle for Moral Realism", *Journal of Philosophy*, 108 (2011): 111–137, and *Moral Knowledge*, Alison Hills, "Moral Testimony and Moral Epistemology", *Ethics*, 120 (2009): 94–127, and Amber Riaz, 2021: "How to Identify Moral Experts", *Journal of Ethics*, 25 (2021): 123–136. On moral understanding see also Amber Riaz, "Moral Understanding and Knowledge", *Philosophical Studies*, 172 (2015): 113–128 and Paulina Sliwa, "Moral Understanding as Knowing Right from Wrong", *Ethics*, 127 (2017): 521–552. Sliwa, "Understanding and Knowing", *Proceedings of the Aristotelian Society*, 115 (2015): 57–74 argues convincingly that understanding why P is just knowing why P.