

## Moral Disagreement and Inexcusable Irrationality

Ralph Wedgwood

Suppose that the relativist view of moral thought is mistaken, and that a broadly absolutist view is correct. Then moral disagreement necessarily involves *error* – in this sense of an incorrect or mistaken belief, a belief in a false proposition – on at least one side or the other. Nonetheless, many philosophers may think that quite often, such errors are entirely *faultless* – at least in the sense that they need involve no irrationality or unjustified belief on either side. After all, it is a familiar fact that perfectly rational or justified beliefs can be false; and when we have such a rational or justified false belief, this typically constitutes an *excuse* for any wrongful acts or decisions that result from it.

As I shall argue here, however, it is not quite so easy to make sense of the idea of faultless moral errors. In fact, I shall argue that completely faultless errors about *ultimate* moral principles are impossible. In a sense that I shall try to explain, mistaken beliefs about ultimate moral principles are in fact *inexcusable*. However, I shall suggest that this conclusion is easier to live with than it first appears.

This idea that having a false belief about ultimate moral principles is no excuse is familiar from the history of philosophy. In a well-known passage in the *Nicomachean Ethics* (1110b29-1111a3), Aristotle seems to argue that even *ignorance* about the universal ultimate principles is no excuse, and that the only kind of error or ignorance that can excuse is ignorance concerning *particular matters of fact*:

Certainly every vicious person is ignorant of the actions he must do or avoid, and this sort of error makes people unjust, and in general bad. This ignorance of what is beneficial is not taken to make actions involuntary. For the cause of involuntary action is not [this] ignorance in the decision, which causes vice; it is not [in other words] ignorance of the universal, since that is a cause for blame. Rather, the cause is ignorance of the particulars which the action consists in and is concerned with, since these allow both pity and pardon.

The idea that being mistaken about ultimate moral principles is no excuse is clearly similar to the famous legal doctrine that ignorance of the law is no excuse (*ignorantia juris non excusat*). This doctrine also seems to be endorsed by Aristotle (1113b30f.):

[Legislators] also impose corrective treatment on someone who [does a vicious action] in ignorance of some provision of law that he is required to know and that is not hard [to know].

I shall argue for a version of this sort of view here. Admittedly, there are two kinds of ignorance: in cases of the first kind, ignorance of a truth involves having no attitude whatsoever (whether belief, disbelief, suspension of judgment, or any other attitude) towards that truth; in cases of the second kind, ignorance of a truth involves having some attitude towards that truth – but an attitude that does *not* involve being completely confident of that truth. It is only the second kind of ignorance that counts as inexcusable in my view; the first kind of ignorance, as I shall explain,

need not involve any genuine error at all. Still, I shall argue that any attitude towards what is in fact a true ultimate moral principle that falls short of complete confidence involves inexcusable irrationality.

More recently, however, a number of philosophers have explored different – and in a sense more *forgiving* – views of moral uncertainty, error, and ignorance. This forgiving view of moral ignorance was pioneered by philosophers like Graham Oddie (1995) and Michael Zimmerman (1997), and developed further by a few other philosophers such as Ted Lockhart (2000), Jacob Ross (2006), and Andrew Sepielli (2009). This recent literature contains several insights, which I shall try to accommodate in the position that I shall advocate here. Nonetheless, fundamentally my position belongs to the opposing camp – which has recently been defended by such philosophers as Elizabeth Harman (2011), Brian Weatherson (2013), and Brian Hedden (forthcoming). According to my opposing position, error and even uncertainty about the true fundamental ethical principles is irrational and fundamentally inexcusable.

### 1. **Ultimate moral principles are *a priori***

Some moral errors are entirely due to errors about non-moral matters. We may set these errors aside. Since it is clearly possible for empirical beliefs about non-moral matters to be false while being at the same time entirely rational and blameless, moral errors that are entirely due to false empirical beliefs about non-moral matters can clearly be faultless. One might be tempted to conclude that any remaining errors must be due to errors about *ultimate* moral principles. I shall later argue that this conclusion does not follow. But for the time being let us focus on cases that do involve such errors about ultimate moral principles.

The main point that I shall argue for in this section is that there are reasons for thinking that each of these ultimate moral principles is in a sense *a priori*. That is, each of these principles is a proposition that – regardless of the evidence that they happen to have – every thinker who possesses the requisite concepts has propositional justification for believing. This is not to say that every thinker is in any strong sense “in a position” to avail themselves of this propositional justification, and to believe the proposition in a doxastically justified manner; it may not be at all feasible for many thinkers to avail themselves of this propositional justification. Nonetheless, the justification for the thinker to believe the proposition is still present, so long as the thinker even possesses the concepts in question.

Perhaps the most important argument for the claim that fundamental moral principles are *a priori* is due to Kant. As Kant says in the Preface to the *Groundwork for the Metaphysics of Morals* (1785, 389):<sup>1</sup>

Everyone must admit that if a law is to hold morally, that is, as a ground of an obligation, it must carry with it absolute necessity: it is not as if the commandment *Thou shalt not lie* holds only for human beings, while other rational beings did not have to take account of it; and the same goes for all other genuine moral laws. Consequently, the ground of the

---

<sup>1</sup> For other passages where Kant insists that the fundamental moral principles must be *a priori*, see Kant (1785, 406-12, 425-27).

obligation must not be sought in the nature of the human being, or in the circumstances in the world in which he is placed, but *a priori* purely in concepts of pure reason.

Kant's point here is that fundamental moral laws apply to all possible rational beings — not just to human beings, but also to angels, Martians and other extraterrestrials, and so on.<sup>2</sup> In this sense, they are as he says “absolutely necessary”. But how could empirical information, about how things actually happen to be with human beings, even be relevant to establishing a necessary principle that concerns all possible rational beings, in absolutely all possible worlds?<sup>3</sup> If moral laws cannot be established on the basis of such empirical information, then it seems that they must be *a priori*.

As I shall understand it, Kant's argument should not be interpreted as implying that *a priori* principles can *never* be known empirically. Even if a mathematical truth can be known *a priori* (as when one knows it directly on the basis of a mathematical proof), perhaps it can also be known empirically (for example, on the basis of the testimony of a credible expert). But in these cases, when an *a priori* truth is known empirically, this seems to be the case only because the empirical basis on which the knowledge is based in some way reflects the existence of a *a priori* justification. Thus, typically, one knows a mathematical proposition *p* on the basis of testimony only because the testimony information derives ultimately from someone who knew *p*, not on the basis of testimony, but on the basis of an *a priori* proof. As I shall interpret Kant's argument, then, such truths can be known empirically – but only in a way that is derivative from a more basic *a priori* justification that these truths have.

As it stands, however, Kant's argument is open to at least two objections. First, even if a truth is, as Kant puts it, “absolutely necessary”, why could it not be the case that the ultimate explanation of this truth involves contingent facts about the actual world – with the result that this truth is knowable only empirically? Secondly, even if the truth in question does not depend for its explanation on contingent facts in this way, such truths could still be empirical. For example, the proposition that water is H<sub>2</sub>O does not depend for its explanation on contingent facts about how things are in the actual world; on the contrary, this proposition is simply a fundamental truth about the constitutive essence or nature of water. Nonetheless, given the concepts that are involved in this proposition, it is clearly deeply empirical.

However, Kant's argument can be amended so that it is also immune to these objections. As I shall try to show, ultimate normative principles do not depend for their explanation on any contingent fact about how things happen to be in the actual world. These ultimate normative principles are just basic necessary truths, which neither require nor admit of any further explanation. As I shall put it, these ultimate normative principles are “fundamental necessary truths”.

Admittedly, as we have seen, there are some fundamental necessary truths of this kind – like the proposition that water is H<sub>2</sub>O – that are undeniably empirical. However, these empirical

---

<sup>2</sup> For a related argument, in favour of the claim that there must be fundamental normative principles that do not depend on any contingent non-normative facts, see G. A. Cohen (2003).

<sup>3</sup> Clearly, Kant is relying here on the principle that he famously articulated in the Introduction to the *Critique of Pure Reason* (1787, 3): “Experience teaches us, to be sure, that something is thus and so, but not that it could not be otherwise”.

necessary truths all seem to involve concepts like ‘water’, the reference of which does not just depend on the concepts’ internal conceptual role in thinking and reasoning, but also depends on further contingent facts about the world (such as the causal relations between one’s use of the concept and the various natural kinds that are actually instantiated in one’s environment). There are other concepts, however, such as logical, mathematical and psychological concepts, the reference of which depends purely on the concepts’ inner conceptual role, and not on any further contingent facts about the world. Fundamental necessary truths that involve only concepts of this second kind seem to be purely a priori.<sup>4</sup>

According to the accounts that seem most plausible, *normative* concepts, the concepts that can be expressed by terms like ‘ought’, ‘good’ and ‘right’ and the like, are also concepts of this second kind: their reference or semantic value depends purely on their inner conceptual role, not on any contingent facts about the relations between the thinker and his environment. So it seems that any fundamental necessary truths that involve only normative, psychological, logical and mathematical concepts will also be a priori, and it will be possible to know those truths on an empirical basis only if that empirical basis reflects the existence of some a priori justification for believing those truths.

In fact, it seems plausible that the ultimate normative principles are fundamental necessary truths of just this kind. First, it seems plausible that normative truths *strongly supervene* on non-normative truths. If one agent  $A_1$  ought to act in a certain way  $W$  at a certain time  $t_1$ , for example, then it is *impossible* for there to be an agent  $A_2$  and time  $t_2$  such that  $A_2$ ’s situation at  $t_2$  is exactly like  $A_1$ ’s situation at  $t_1$  in every non-normative respect, without its also being the case that  $A_2$  also ought to act in the same way  $W$  at  $t_2$ . It follows that there is some non-normative relation between an agent and a time such that it is *necessary* that whenever an agent stands in that non-normative relation to a time  $t$ , the agent ought to act in the way  $W$  at  $t$ .

In general, it follows from this idea of the strong supervenience of the normative on the non-normative that whenever a sequence of items exemplifies a normative relation, it also exemplifies a non-normative relation that necessitates that normative relation.<sup>5</sup> Every sequence of items that exemplifies a normative relation exemplifies some non-normative relation such that it is a necessary truth that every sequence of items exemplifying that non-normative relation also exemplifies that normative relation. Some of these necessary truths can be explained on the basis of other such truths. But these explanations will have to start somewhere, and so there must be some ultimate necessary truths of this sort that do not have any further explanation at all.

It also seems plausible that normative truths strongly supervene on truths that can be stated using

---

<sup>4</sup> For an argument for this point, see Bealer (1987) on the limits to what he calls “scientific essentialism”. It would take too long to offer an explanation here of exactly *why* necessary modal truths involving only concepts of this sort must be knowable a priori. Very roughly, I would suggest something along the following lines: if such a truth holds it does so as a result of the essence of the objects, properties or relations that are mentioned in that truth; if the inner conceptual role of a concept determines the concept’s semantic value, then it must also determine the essence of that semantic value; and if that is the case, then the nature of the concept must provide a priori justification for believing the truth about the essence of that semantic value.

<sup>5</sup> See Wedgwood (1999, 201–02) for a slightly fuller argument for this point.

non-normative concepts the reference of which does not depend on contingent facts about the thinker's environment. It seems highly doubtful whether the fact that one's situation involves water (rather than some other substance whose mental effects are just like those of water) could make any difference to what one ought to do, unless it also made a difference that could be described in psychological or sociological terms — such as making a difference to what type of act will count in the situation as keeping a promise or speaking truthfully or the like.

More specifically, it may be that what an agent ought to do at a time  $t$  strongly supervenes on *mental* and *social* facts about the agent's situation at  $t$ . So there are fundamental necessary truths to the effect that in all possible worlds, whenever such-and-such a mental and social relation holds between an agent and a time, the agent ought to do something of such-and-such a kind. In general, whenever a sequence of items exemplifies a normative relation, that sequence of items also exemplifies some (possibly highly complex) mental and social relation such that it is a necessary truth that every sequence of items exemplifying that mental and social relation also exemplifies this normative relation. As we have seen, some of these necessary truths will be fundamental necessary truths that require no further explanation; these fundamental truths will count as ultimate normative principles.

As I have already mentioned, it also seems plausible that such mental and social relations can be specified using concepts (like psychological, logical, and mathematical concepts) the reference of which is determined purely by the concepts' inner conceptual role. As I have also suggested, it seems that necessary modal truths involving only concepts of this sort must be a priori. As I have argued, it seems plausible that the ultimate normative principles are necessary modal truths that can be specified using only concepts of this sort. This then gives us a reason to be sympathetic to the Kantian position that the ultimate normative principles are a priori.

## 2. Probabilistic accounts of rationality

Such necessary *a priori* truths are true in all possible worlds. Moreover, these truths are not only true at all metaphysically possible worlds, like the proposition that *water is H<sub>2</sub>O*. It seems that these truths are necessary in a different sense as well. They are like logical truths and analytic truths in being *epistemically necessary*. In fact, what seems characteristic of epistemically necessary truths is precisely that they are a priori in this way: every thinker possessing the concepts that are necessary for even considering or entertaining these truths – whatever empirical evidence those thinkers may have – has justification for believing these truths. This seems to be the best account of the kind of necessity that is shared by logical truths and analytic or conceptual truths – including even those unusual analytic truths that are metaphysically contingent (like truths of the form 'If  $p$ , then actually  $p$ ', or Saul Kripke's famous example concerning the length of the meter rod in Paris).

These epistemically necessary truths will play an important role in any account of rationality. To fix ideas, I shall explore the role that these truths play in the kind of account of rationality that I favour – namely, a broadly *probabilistic* account.

According to probabilistic accounts of rationality, for every thinker and every time, there is a *probability function* that rationally should be guiding the thinker at that time. (This approach could be refined by allowing that there may not always be a *unique* probability function that

rationally should be guiding the thinker at this time, but sometimes there is only a *convex set of probability functions* that rationally should be guiding the thinker at the time, or the like. But these details will not matter for our purposes.) Just to have a label, let us call this probability function the *rational probability* function for the thinker at the time.

Every probability function, we may assume, is defined over a space of possible worlds. The rationality probability function is not defined over a space of metaphysically possible worlds, since it seems that a perfectly rational thinker could be guided by a probability function that assigns non-zero probability to metaphysically impossible propositions like the proposition that *water is not H<sub>2</sub>O*. However, it does seem plausible that the rational probability function must assign probability 1 to all logical and analytic truths; and the best account of why this is the case, I propose, is that this probability function is defined over a space of *epistemically possible worlds*.

Since as I have argued, every a priori truth is epistemically necessary, every such truth must have probability 1 in any probability function that is defined over these epistemically possible worlds. Since the rational probability function is defined over these worlds, the a priori truths must have probability 1 according to the rational probability function. Since, as I have also argued, ultimate moral principles are a priori, these ultimate moral principles must have probability 1 according to the rational probability function.

The way in which this rational probability rationally should “guide” the thinker presumably involves the following two points:

- If the thinker has any sort of credence in a proposition *p* at a time, that credence should match the probability of *p* according to the rational probability function for the thinker at that time.
- The thinker’s choices and intentions at each time should maximize some kind of expected value that is defined in terms of the rational probability function for the thinker at that time.

For our purposes, the most relevant point is the first. If a thinker is to be perfectly rational, and the thinker has any sort of credence in an a priori truth, then that credence must be complete maximum confidence – the level of credence that matches probability 1.

Given a probabilistic conception of rationality, it follows that the only perfectly rational attitude to have towards these a priori truths is to have maximum confidence in them. This approach allows that it can be rational to have no attitude whatsoever towards these a priori truths, but if one does have an attitude, that attitude must be maximum confidence.

Even suspension of judgment is arguably irrational, since it too is some kind of broadly doxastic attitude that falls short of complete certainty. It follows that every perfectly rational person must never assign any credence other than 1 to such propositions. It need not be irrational in any way if one simply never considers such a proposition at all. But to be perfectly rational, anyone who does consider such a proposition must have no doubt about them at all.

It is widely recognized that this is what probabilistic accounts of rationality would have to say about logical truths. (Indeed, many proponents of probabilistic approaches to rationality unnecessarily impose the stronger requirement of *logical omniscience* – not only *forbidding* any

attitude *other* than credence 1 in any logical truth, but positively *requiring* credence 1 in all logical truths.) This is often recognized to be a strong and initially surprising consequence of the probabilistic approach. (After all, some logical truths are fantastically difficult to prove: so how can it be that it is irrational to have any attitude towards these truths that falls short of maximum confidence?)

However, many philosophers also think that the probabilistic approach to rationality has proven its power in epistemology and the philosophy of mind, by explaining a wide variety of phenomena that need to be explained – while none of the alternatives to probabilistic approaches have had comparable explanatory success. So many philosophers think that we should find a way of reconciling ourselves to the idea that the only perfectly rational attitude to have towards any logical truth is the attitude of maximum confidence.

My main point in this section is just to propose that these probabilistic accounts should also say the same point about *all* a priori truths – including (if I am right) ultimate moral principles as well.

### 3. Subjective rightness and excuses

If it is to a certain degree irrational for one to have a certain attitude, then – at least so long as it is possible (in the relevant sense of ‘possible’, whatever that may be) for one to have an attitude that is less irrational than that – it seems that one is *rationally required* not to have that attitude. In some sense, one *ought not* to have that attitude.

If one has an attitude that one ought not to have, doesn’t it follow that one merits a kind of *criticism* – in some sense of the term, even a kind of blame? But surely – it might be thought – we are not always blameworthy whenever we fail to have an attitude of maximum confidence towards a logical truth, even if it is possible to have an attitude of maximum confidence at the relevant time. Such failures – it might be thought – can surely be entirely blameless.

A case in which one is blameless for thinking in a certain way in spite of its being true that one *ought not* to think in that way is a case of an *excuse*. In this section, I shall explore the suggestion that there are cases in which we think in a way in which we are rationally required not to think and yet are entirely blameless for thinking in that way. In other words, I shall explore the suggestion that it is possible for violations of a rational requirement to be *excusable*. I shall argue that this is not in fact possible. All avoidable irrationality is inexcusable.

The paradigmatic example of an excuse is a case in which one acts or thinks in a way that is objectively wrong but subjectively right, on account of one’s having a rational false belief. For example, you kill a person because you rationally but falsely believe that the person is tempting to kill you, and that killing the person is the only way for you to defend yourself against his attack. In this case, your act is objectively wrong – no one who knew the truth could rightly assist you in your killing this innocent person – but it is blameless, excused because of your rational false belief.

This sort of excuse clearly does not apply in the cases that we are concerned with, because they are *ex hypothesi* cases of irrational thinking. Are there any other varieties of excuse that might apply in these cases?

In fact, it is not clear that there are any other kinds of excuse. Many of the factors commonly classified as excuses seem on closer inspection to belong in other categories. For example, it is often thought that coercion and duress are excuses. But, in the strict sense that we are using the term here, you have an excuse only for acting in a way in which you *ought not* to have acted (but for which you are blameless, on account of the excuse). In the paradigmatic cases of coercion and duress, you in fact act precisely as you in your unfortunate situation *ought* to act. The bank clerk who opens the safe and hands the money to the bank robber is acting precisely as she ought to act in her extremely unfortunate situation. She may well deserve praise – at least to a modest degree – for responding in this prudent fashion, rather than foolishly attempting to resist the bank robber. It would be clearly permissible for you to help her to open the safe and give the money to the bank robber, even if you knew all the relevant truths about the situation. In short, these cases of coercion and duress involve a *justification* rather than an excuse.

Another alleged excuse is *incapacity*. But if the agent is incapable of acting in a certain way, then given the plausible principle that ‘ought’ implies ‘can’, it cannot be true that she ought to act in that way. Incapacity removes the agent from the scope of the ‘ought’. So incapacity also cannot be an excuse – it cannot lead to cases in which is blameless for acting or thinking in a certain way in spite of the fact that she ought not to have acted or thought in that way.

Other factors, such as provocation and emotional distress and the like, seem to function by increasing the *difficulty* of acting or thinking as one ought to – but without (like incapacity) making it strictly *impossible* for one to avoid acting or thinking in the relevant way. However, difficulty comes in degrees, ranging all the way from the kind of extreme difficulty that verges on sheer impossibility, to mild awkwardness and inconvenience. Since it is clear that the milder levels of difficulty do not ground total excuses, it seems that difficulty should be interpreted, not as grounding excuses, but as a *mitigating* factor, which somehow *lessens* the degree of wrongness or the degree of blameworthiness without removing it altogether.

In short, it seems that if one was capable of a greater degree of rationality than one actually exemplified, it follows that one thought in a way in which one ought to have, and one has no excuse. In short, all avoidable irrationality is inexcusable.

In the cases that currently concern us, one has an attitude towards a true ultimate moral principle, but one’s attitude falls short of maximum confidence. I have argued that this is at least to some degree irrational. If it was possible for one to have maximum confidence in that principle, then one is violating a rational requirement – which as I have argued, is strictly inexcusable.

#### 4. Mitigating factors

The previous section in fact points us towards the best way of reconciling ourselves with the seemingly extreme position that we have arrived at. Mitigating factors may in fact be ubiquitous. Even if we do always merit criticism for thinking in ways that are irrational (at least if it was possible for us to think in less irrational ways), the degree of criticism that we merit may be severely mitigated, for two reasons:

- First, we might only manifest a very low degree of irrationality (or only a degree of rationality that falls very slightly short of the greatest degree of rationality of which we are capable).

- Secondly, it may also be extremely difficult – though not quite impossible – for us to achieve higher levels of rationality.

Together these two points imply that even if we are for much of the time at least somewhat less rational than we could have been, the degree of blame or criticism that we merit may be so massively mitigated that it is effectively imperceptible. When the degree of criticism that we merit is so low, it might typically be wrong to express such criticism in conversation, and even pointless to focus on the fact that one merits this degree of criticism in one's own private thoughts. This is still compatible with its being true that strictly speaking one does merit a very low level of criticism.

In general, we may find it hard to distinguish between cases in which we do not merit criticism at all and cases in which we merit an extremely low level of criticism. (Mark Schroeder makes a similar point about how it may be difficult to distinguish between cases where we have no reason at all for a certain course of action and cases where we have a very weak or lightweight reason for that course of action.) So the claim that having an attitude other than maximum confidence in an ultimate moral principle (when it was possible for to have such an attitude of certainty) is a case of inexcusable irrationality is in fact less extreme and implausible than it initially sounds.

In addition to this point about the ubiquity of mitigating factors, there may also be a second way in which we can reconcile ourselves with the seemingly extreme position that we have arrived at here. I have argued here that *ultimate moral principles* are a priori. (Presumably, the logical consequences of such ultimate moral principles are also a priori.) However, it is not clear that any actual human being has ever had any attitudes at all towards these a priori moral principles. Admittedly, a few philosophers think that these ultimate principles are quite simple – such as the classical utilitarian principle, to give one example. But most philosophers suspect that in fact these ultimate principles would be fantastically complicated – so complicated that they may in fact never have occurred to any actual human being.

If we in fact have no attitudes at all towards these ultimate moral principles (or their logical consequences), then nothing that I have said implies that we are being irrational in any way. According to the version of the probabilistic approach that I am assuming here, it is irrational to have attitudes other than maximum confidence in a priori truths, but it is not irrational to have no attitudes at all.

Now, as Daniel Star (2015) has persuasively argued, we ordinary human beings do have basic knowledge of certain moral principles. Our knowledge of such principles is “basic” in the sense that our justification for believing them does not depend on our having inferred them from any other moral propositions. But the moral principles of which we ordinary human beings have this sort of basic knowledge are such mid-level principles as that (for example) *it is normally wrong to hit people and to steal things*. These mid-level principles are not ultimate principles that provide the fundamental explanation of the truth about morality. They depend for their truth on the ultimate moral principles together with a host of background empirical assumptions about such matters as the normal consequences of hitting people and of stealing things. Since we do not know the ultimate moral principles, we also do not know how exactly the mid-level principles that we do know depend on this mixture of empirical fact and ultimate moral principles.

Since these mid-level principles are empirical but also non-inferentially justified, we cannot

assume that whenever there is a moral disagreement that is not entirely due to an empirical disagreement about non-moral matters, the disagreement must be due to a difference over the ultimate moral principles instead. The disagreement may fundamentally concern these mid-level principles.

These mid-level principles are clearly contingent. It also seems that they are not a priori: even though they are not believed as a result of any inference from prior beliefs, their justification can still be defeated if new empirical information comes to light. In that sense, these beliefs are empirical. A fortiori, our moral beliefs about particular events (like my belief that it was wrong of Julius Caesar to cross the Rubicon) are also empirical. In short, it seems that the overwhelming majority of the moral beliefs that we have are empirical. The arguments that I have given here do not raise any problems for the thesis that false beliefs about empirical matters can be perfectly rational and so entirely excusable. So the issue that I have been discussing here in fact concerns only a vanishingly tiny proportion of our moral beliefs – if indeed it concerns any of our actual beliefs at all.

## 5. Back to disagreement

Assuming with the anti-relativists that at least one side in every genuine disagreement must have a false belief, I have argued that there cannot be a disagreement about ultimate moral principles in which both sides of the disagreement are perfectly rational.

Instead, in real-life cases, each side in the disagreement should recognize that they most likely neither of them perfectly rational, and so each side in the disagreement merits at least some mild criticism for the irrational aspects of their thinking. Nonetheless, even the steps that would lead each side towards more rational thinking will very often involve each side's regarding it as more likely that they are right while the other side is wrong. This, I suggest, is the kind of "reasonable disagreement" that we have to deal with in political life. Indeed, it seems plausible to me that we have such "reasonable disagreements" not just about ultimate moral questions, but also about the mid-level moral principles that (according to views like Daniel Star's) many ordinary human beings know. In fact, I would say that we have such "reasonable disagreements" not just about ethical and philosophical issues, but about the difficult questions of economics and social theory as well. All that I have argued here is that this politically important kind of "reasonable disagreement" is not strictly speaking a matter of perfectly rational disagreements about ultimate moral principles.

## References

- Bealer, George (1987). The philosophical limits of scientific essentialism. *Philosophical Perspectives* 1:289-365.
- Cohen, G. A. (2003). "Facts and Principles", *Philosophy and Public Affairs* 31 (3): 211–45.
- Guerrero, Alexander A. (2007). "Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution." *Philosophical Studies* 136: 59–97.

- Harman, Elizabeth (2011). Does moral ignorance exculpate? *Ratio* 24 (4):443-468.
- Hedden, Brian (forthcoming). "Does MITE Make Right? Decision-Making Under Normative Uncertainty", *Oxford Studies in Metaethics*.
- Kant, Immanuel (1785). *Groundwork of the Metaphysics of Morals*. Cited by the pagination of the Royal Prussian (later German) Academy edition (Berlin: 1900–), Vol. IV.
- Kant, Immanuel (1787). *Critique of Pure Reason*, 2<sup>nd</sup> (B) edition (Riga: Hartknoch).
- Lockhart, Ted (2000). *Moral Uncertainty and Its Consequences*. (Oxford: Oxford University Press.)
- Oddie, Graham (1995). "Moral Uncertainty and Human Embryo Experimentation." In *Medicine and Moral Reasoning*, ed. K. W. M. Fulford, Grant Gillett, and Janet Martin Soskice (Oxford University Press).
- Rosen, Gideon (2003). "Culpability and ignorance", *Proceedings of the Aristotelian Society* 103 (1): 61-84.
- Ross, Jacob (2006). "Rejecting Ethical Deflationism." *Ethics* 116: 742–68.
- Sepielli, Andrew (2009). "What to Do When You Don't Know What To Do." In *Oxford Studies in Metaethics*, ed. Russ Shafer-Landau (Oxford: Oxford University Press).
- Star, Daniel (2015). *Knowing Better* (Oxford: Oxford University Press).
- Weatherson, Brian (2013). "Running Risks Morally", *Philosophical Studies* 167: 1-23.
- Wedgwood, Ralph (1999). "The price of non-reductive moral realism". *Ethical Theory and Moral Practice* 2 (3): 199-215.
- Zimmerman, M. (1997). "Moral responsibility and ignorance", *Ethics* 107 (3): 410-426.