THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Obvious Objections

Citation for published version:

Ridge, M & McKeever, S 2015, Obvious Objections. in *Intuition, Theory and anti-Theory*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198713227.001.0001>

Digital Object Identifier (DOI):

[10.1093/acprof:oso/9780198713227.001.0001](https://doi.org/10.1093/acprof:oso/9780198713227.001.0001)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Intuition, Theory and anti-Theory

Publisher Rights Statement:

(c) Copyright Oxford University Press, 2015. All Rights Reserved.

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

University Press Scholarship Online

Oxford Scholarship Online

Intuition, Theory, and Anti-Theory in Ethics

Sophie Grace Chappell

Print publication date: 2015

Print ISBN-13: 9780198713227

Published to Oxford Scholarship Online: April 2015

DOI: 10.1093/acprof:oso/9780198713227.001.0001

Obvious Objections

Sean McKeever

Michael Ridge

DOI:10.1093/acprof:oso/9780198713227.003.0003

[–] Abstract and Keywords

People often say that some normative ethical theory, act consequentialism perhaps, faces some objection based on ‘obvious intuitions’. It is also sometimes said—e.g. by Bernard Williams—that a theory like act consequentialism can make some things *too* obvious; can make obvious things that *aren’t* obvious but deeply morally problematic and troubling even if correct. Ridge and McKeever agree that objections of Williams’ sort can sometimes have some weight; but how much weight, and how often? They remind us that the anti-theorist’s deeper aim is often to insist on the subtlety and variety of the moral landscape. But, the more subtle and variable the moral landscape can be, the less sure we can be that we know which cases truly are the ‘hard ones’. And if we cannot be confident which cases are hard, then we cannot have strong reason to reject a theory for miscategorizing hard cases.

Keywords: intuitions, normative ethical theory, obvious, act consequentialism, Bernard Williams

I

Perplexity, it might be thought, is the philosopher's mortal enemy, always to be vanquished, never accommodated. For normative theorists, the typical weapon of choice is an explanatory theory couched in general or principled terms, one that lays out the various considerations that can make acts right and wrong and how these considerations balance against one another. Though open to important variations, there is a familiar choreography to the ensuing battles. We begin with an array of relatively specific claims in which we place great confidence; call these intuitions. The theory, it is hoped, will more or less preserve these intuitions, forcing minor revisions here, shedding a bit of light there, but mostly leaving things as they stood. After all, forcing revision to firmly held beliefs introduces perplexity of its own. How could we have ever been so sure of what we now reject?! Some significant revision is no doubt inevitable, but it is regrettable all the same, and if it gives rise to additional perplexity, then this is just another philosophical skirmish to be fought.

But, of course, our intuitions cover only small and irregular patches of the moral terrain. Intuitions are not (at least as we use the term here) to be identified with whatever a subject might be brought to say about a case if pressed. Whether they are pre-reflective is not essential, but a depth of confidence is. Intuitions are claims we place confidence in and in which we think ourselves right to have confidence (even if we cannot say why). Thus understood, it should be clear that intuitions are the exception, not the rule of moral life. For there are innumerable issues on which we are uncertain, confused, conflicted, ambivalent, or schizophrenic. And then there are the questions which have engaged us so little that we hardly have thoughts about them at all. Much of this perplexity is surely due to uncertainty about underlying non-moral facts—facts that may be morally relevant but that can be characterized in a non-moral way. But we will assume here that this **(p.41)** is not always the case. Much of our perplexity is at bottom moral; it is perplexity about what sorts of considerations are morally relevant and how they balance against one another. Normative theories, it is commonly thought, should yield answers about these perplexing matters. The point of normative theorizing is not merely to build a foundation that will secure our pre-existing views, but to bring light to previously dark corners.

One might worry whether the project just sketched can succeed at all, but here we ask a different question: can it succeed too well? The basic idea is this. Sometimes perplexity is in order; a moral question is such that we should be uncertain or conflicted about its answer. If this is right, then perhaps theories that yield clear and straightforward answers to 'hard' cases are, on that ground alone, flawed. If the theory renders obvious what should not be obvious, can this be a mark against the theory?

Bernard Williams lodged a particularly well-known version of this complaint in his argument against utilitarianism. After introducing his now famous examples—the first involving a chemist (George) who is offered a job working on the production of chemical weapons (something he abhors), and the second involving a traveller (Jim) who is enticed

by a warlord to save nineteen captured Indians but only by shooting one himself—Williams remarks:

To these dilemmas, it seems to me that utilitarianism replies, in the first case that George should accept the job, and in the second, that Jim should kill the Indian. Not only does utilitarianism give these answers but, if the situations are essentially as described and there are no other special factors, it regards them, it seems to me, as *obviously* the right answers. But many of us would certainly wonder whether in (1), that could possibly be the right answer at all; and in the case of (2), even one who came to think that perhaps that was the answer, might well wonder whether it was obviously the answer. (Williams, 1973: 99, emphasis in original)

As Williams goes on to make clear, these examples are not intended as standard counterexamples. Williams is willing to grant that George should take the job and Jim should shoot an Indian. The complaint is that utilitarianism is making obvious something that isn't or shouldn't be. Famously, Williams goes on to proffer an explanation for why utilitarianism renders obvious what should not be. Utilitarianism cannot, Williams claims, make sense of the value of integrity which is crucially at play in the examples. This further explanation (left unexamined here) is meant to bolster the case against utilitarianism. But the argument at issue here is not initiated by this further explanation. If anything, matters are reversed. It is only because we think it a flaw that the theory makes obvious what shouldn't be that we look for a deeper explanation of how and why the theory is oversimplifying matters.

(p.42) Lest Williams's use of this form of argument seem idiosyncratic, here is another example, this one taken from D. D. Raphael's 'The Standard of Morals' (1975):

Sartre gives us a good instance of moral conflict in his young friend during the Second World War who had to choose whether to go off and join the Resistance or to stay with his widowed mother. The advantage of this well-known example is that it is taken from real life and might have faced any of us. If you were to think of the young man's dilemma simply in terms of treating persons as ends and of the number of persons affected, his proper choice would be obvious; but in fact, of course, the right choice is not in the least obvious. (Raphael, 1974–5: 9–10)

Although the argument here is enthymematic, it is clear that the target is Kantian moral theory. So whereas Williams used this form of argument to criticize utilitarianism, Raphael uses it instead to criticize utilitarianism's main rival, Kantianism.

Our interest, we should stress, is not utilitarianism in particular, or even an evaluation of Williams's own argument, which has already received ample discussion. Instead, our interest is in the form of arguments along these lines, and whether they ever can be sound. Our qualified and cautious answer is yes, though such arguments are liable to be very weak. Perhaps surprisingly, one ground for such weakness derives precisely from a feature of morality that is often emphasized by those who are sceptical of general normative theory. But we will come to that conclusion in a roundabout way by first

considering a variety of ways that such arguments can go awry.

II

In an oft-repeated quip some philosopher insists that a point is obvious, whereupon another retorts, 'Yes, that's right, but the trouble is that you think your point is obviously true while others of us think it obviously false.' Fair enough (sometimes); philosophers can be tendentious. But an on-target quip does not a general theorem make. From the fact that appeals to obviousness are often lame attempts to scuttle disagreement, it does not follow that we are never in substantial agreement that some point is true and obviously so. Were that so, it would present at least a dialectical barrier to arguments of the sort we are considering. But on reflection, we do often find ourselves in substantial agreement that some claim is true and obvious.

We offer no definition of substantial agreement, but it is a looser standard than consensus. It may be true that there is no claim which we unanimously agree is both true and obvious, but it would be churlish to insist on unanimity. In the present discussion, we set to the side the serious challenges to moral 'intuitions' (p.43) and the prominent role they play in constructing and assessing normative theories. Arguments such as Williams's are not supposed to depend upon general scepticism about moral intuitions. At issue instead is whether, given that one measures a theory in part by its ability to fit and explain our intuitions, a theory can go wrong by, somehow, conflicting with our intuitions about what is obvious. Against that backdrop, it cannot be assumed that all appeals to what is obvious are tendentious because there is no agreement about what is obvious. By the same token, there is virtually never consensus about moral intuitions, but there may be widespread agreement. And just so, there may be (sometimes) widespread agreement about what is obvious. The appeal to intuitions about the obvious may yet prove tendentious. Our present point is only that this conclusion cannot be established by quip or by the casual and false assumption that there is insufficient interpersonal agreement about what is obvious.

Turning to more worrying problems, let *p* stand for a particular moral claim and *T* for a normative theory. Here is one form that arguments such as this might take:

- (i) *p* is not obvious (even if true).
- (ii) If *T* were true, *p* would be obvious.

So,

- (iii) *T* is false.

As others have already rightly argued, this argument form is not promising. In a given case, premise (ii) is almost surely false, and any illusion of its truth is likely to be the product of confusion. The truth of a *theory* just does not typically confer obviousness on its *implications*. One reason for this, noted by Fred Feldman in a discussion of Williams, is that a true theory can itself be quite unobvious, even when it is known to be true. The truth of utilitarianism can hardly establish that Jim killing an Indian is, morally speaking, a

'no-brainer.' This is due in part to what we all know: utilitarianism itself is *not* a no-brainer (Feldman, 1986; see also Lenman, 2004).

If (ii) seems plausible, this may be on account of confusing it with what is more plausibly true. It is obvious, it may seem, that *T implies p*. In the moral case, however, *T* will imply *p* only relative to certain contingent facts. Call these *F*. For example, it is commonly assumed that utilitarianism implies that Jim should shoot an Indian relative to the assumptions that (i) Pedro will honour his pledge to release the other hostages if Jim shoots and that (ii) he will kill all the hostages if Jim refuses. Some are suspicious of the implication even then (Lenman, 2004). Even setting that worry to the side, however, if one replaces *F* with other facts one **(p.44)** may lose the implication. Since the facts on which the implication depends may not be obvious, the implication may not be obvious either. Finally, even if the further-needed facts are true and obvious, the implication itself may not be obvious. What follows from Kant's Formula of Universal Law, for example, is hardly obvious even when knowledge of the underlying facts is safe and secure.

What is plausibly true, then, is only this:

(iv) (*T* implies *p*) given *F*.

Can the argument be rehabilitated on this basis? Let us just suppose that it is obvious that (*T* implies *p*) given *F*. It might then be thought that:

(v) If *T* were obvious and *F* were obvious, then *p* would be obvious.

Arguments relying on (v) would have limited application, of course, since theories and facts may seldom cooperate by being obvious. One might try to turn this into an argument with the following form:

- (1) *p* is not obvious (even if true).
- (2) If *T* is obvious and *F* is obvious, then *p* is obvious [on the grounds that (*T* implies *p*) given *F*].
- (3) If *T* is true, then *T* is obvious.
- (4) *F* is obvious.
- (5) If *T* were true, then *p* would be obvious.
- (6) So, *T* is not true.

Even when the obviousness of initial facts and theory are in place, we really have here only the ingredients of a paradox. Premises (1) to (4) are an inconsistent quartet, but it is an open question which of them we should reject.

If one initially accepted all four claims, then rationality would require one to reject one of them. We suspect, however, that in most cases the natural adjustment would be to reject (3) or (4). Philosophical theories of interest are seldom such that if they are true then they are obviously true, and contingent facts which are relevant in these contexts will also often not be obvious. So these are in a sense the smallest steps needed to restore coherence to our judgements. One could reject (1), but this would require both

changing one's view of p 's obviousness *and* taking a stand on p 's truth. By contrast, we can restore coherence simply by adjusting our view of T 's (or F 's) obviousness (see also Lenman, 2004).

Some critics of Williams's argument are willing to accept (2) (Feldman, 1986). On close inspection, however, the support one might give for (2) is actually very weak, so we could also reasonably abandon (2). Obviousness does not seem to transmit from theory to implications under even these highly restricted **(p.45)** conditions. Sometimes a theory is obvious as is the fact that the theory carries some implication. And yet the implication may remain stubbornly surprising and unobvious. For example, I find the basic rules of arithmetical calculation obvious. Moreover, it is obvious that if these rules are correct, then $2^9 = 512$. But it remains surprising that $2^9 = 512$. The rapidity with which the numbers increase when raised to a higher power can stubbornly confound our expectations. To take another example, one can, with training, find the rules for calculating compound interest quite obvious. And yet the gross amount required to pay off a long-term loan may continue to have the power to startle.

We have so far left unexamined just what it is for some claim to be obvious. Implicit in the discussion, however, has been the idea that the obvious is to be contrasted with what is hard to see or surprising. Thus understood it is perhaps predictable that we have no firm rules for when the implications of some theory will be obvious or not. Whether some claim is obvious says more about us than it need say about what might explain or justify that claim. Even if we start out in substantial agreement about what is obvious, we cannot project that agreement into the future with any confidence.

This brings us to a final and related trouble with the present way of construing the argument form we are considering. Something that is obvious at one time need not be obvious at another time. Theories do, on occasion, have the effect of making obvious for us what was previously uncertain. We can grant that it is not obvious that Jim's moral duty is to shoot an Indian. And even if utilitarianism is true and implies that Jim should shoot an Indian, the unobviousness of utilitarianism is liable to leave the unobviousness of Jim's duty undisturbed. But suppose that we became utterly convinced that utilitarianism is true, and suppose this had the effect of making it obvious to us that Jim's duty is to shoot an Indian. What would be objectionable about this turn of events? We are looking for an argument that does not depend on ordinary counterexample, so we cannot assume that it is *not* Jim's duty to shoot an Indian. If, on the other hand, it is Jim's moral duty to shoot an Indian, then there seems nothing wrong with utilitarianism if this becomes obvious.

This last point must not be confused with the intriguing possibility that there would be something wrong with a person who found it obvious that someone like Jim should shoot an Indian. In the same way that surprising mathematical results can lead us to 'check our work', a surprising moral conclusion might lead us to 'check our deliberation'. Moreover, there may be good utilitarian reasons for being disposed to 'check one's work' before quickly dispatching an innocent Indian. So perhaps a good person would never find the rightness of shooting an Indian obvious. But again, this says more about us than it does

about the truth of utilitarianism.

(p.46) The lesson of the foregoing discussion is this. If obviousness is subjective—simply a matter of how surprising or clear cut things seem to us—then we cannot cast doubt on a theory by claiming that it would make some non-obvious claim obvious. First, it is not generally true that theories have the power to make their implications obvious. Secondly, even if a theory did have this effect, it would say more about us than about the theory.

III

To rehabilitate the argument under consideration, we need to look in another direction and use a more objective conception of what is obvious. In this section, we explore a conception of the obvious/unobvious contrast according to which it is cashed out in terms of simplicity of explanation. When we say that it is not obvious that Jim should shoot an Indian, we may be thinking that this is a ‘hard case’, a case that it does not admit of a simple explanation. By contrast, when we speak of what is obviously morally so, we may be thinking of things whose moral explanation is simple.

We here set aside the difficult task of characterizing simplicity of explanation. We do assume that simplicity of explanation is sufficiently objective even if it is not utterly independent of our particular human capacities and cognitive interests. We assume only that if norms of simplicity are relevant to theory assessment then these norms will be sufficiently stable across persons and times. When a theory explains some phenomenon we can assess that explanation for its simplicity/complexity by applying shared and stable standards. That should be enough for us at least to avoid the pitfalls explored in previous sections.

We can now consider a new version of our target argument:

- (i) p (though it may be true) does not admit of a simple explanation.
- (ii) If T were true, then there would be a simple explanation for p.

So,

- (iii) T is false.

The difficulty with this form of argument lies in the first premise. How can we know that (i) is true?

We might, of course, believe that p requires a non-simple explanation because we have worked out a specific and complex explanation. But how do we know that this is the correct explanation and that no simpler explanation will do? The natural suggestion is that our complex explanation for p is embedded in a broader theory which is itself well justified. This theory must be a rival to T. This is **(p.47)** clearly question-begging, however. If the claim that p cannot be simply explained depends on a prior rejection of T, then that claim cannot be deployed in a persuasive argument against T.

One might try to press this point further by claiming that we have no reliable intuitions about how complex explanations must be except those that are driven by broader judgements about the theories relative to which explanations are given. But this seems to go too far. Sometimes we can have confidence in a particular explanation even without having fitted that explanation into a broader framework that one would call a theory. For example, if one is in a position to save either a group of ten or a (non-overlapping) group of two strangers, most agree that one should save ten, *ceteris paribus*. Furthermore, many would agree that the explanation for this duty is that ten lives is more than two. One can find compelling this simple explanation without first being wedded to a broader theory, such as utilitarianism, into which this explanation might be fitted. Even those with non-consequentialist predilections might hope that they can accommodate the idea that what explains the duty to save ten is a simple fact about greater numbers.

We are here relying on a distinction between the particular facts that explain a particular moral truth and the broader theory which sets out a story about how those facts succeed in being explanatory across a thorough range of cases. This is a distinction that has been much emphasized by particularists such as Jonathan Dancy (2004) who are keen to allow that one can frame sound moral explanations without implicating universal principles. But it is a distinction that can, and should, be accepted even by those not attracted to particularism. If one denies the distinction one is left saying that the 'complete explanation' of any particular moral truth involves a full moral theory. One cost to this move, however, is that it undermines the contrast between simple and complex explanations of moral phenomena.

For example, assume for the moment that some version of utilitarianism is true. Assume further that one should save ten rather than two, *ceteris paribus*. Finally, assume that governments ought to raise revenue via graduated income taxes rather than consumption taxes. This should leave us able, we think, to say that some moral truths admit of simple explanations (such as saving ten rather than two), while others would require complicated explanations (such as the best legal regime for raising revenue). However, if utilitarianism must itself be part of the 'complete explanation' of any true moral claim, then the complete explanation for each of these moral claims must be highly complex. It needn't follow that the two explanations are equally complex, since there remain other differences between the cases. Nonetheless, it seems safe to say there would be no simple moral explanations, at least not by our ordinary baseline of simplicity in such matters.

(p.48) Indeed, we typically understand such simpler explanations to provide common ground between partisans of different moral theories (and opponents of moral theory, for that matter). It is because the consequentialist and the deontologist agree that (e.g.) the fact that one could save more strangers by performing a given action provides a kind of explanation of why one ought to perform the action that they can then have a meaningful debate about which theory best explains why that fact explains the rightness of the action. As Jonathan Dancy has argued, we must be cautious about assuming that simply because one fact (the truth of a given moral theory, say) *enables* some other fact (the fact about

saving a greater number of strangers, say) to explain some phenomena that the enabling fact is (or must be) itself part of any adequate explanation of the phenomena explained.

So we accept that we can frame sound moral explanations without prior recourse to theory. And as the case of saving the ten suggests, sometimes we can be sufficiently confident in the simplicity of an explanation that this should influence our choice of theories. If non-consequentialists cannot show that facts about the greater number can (in suitable cases) explain what is right, then so much the worse for non-consequentialism. The challenge for non-consequentialists is to make sense of this simple explanation without compromising their other theoretical commitments.

But granting all of this is not enough to vindicate the argument form we are considering. For that argument relies on the assumption that some moral claim (if true) *requires* a non-simple explanation. By contrast, we have so far only granted that sometimes we are rightly confident that some claim *admits* of a simple explanation. Put another way, we may sometimes have good reason to reject a theory because it needlessly complicates matters. But it is a far stretch from here to the claim that we can reject a theory for unduly simplifying matters.

Once we focus on this dimension of the argument, moreover, it is doubtful we are entitled to such a premise. First, there is an inductive point. Intellectual progress is regularly marked by the discovery of surprisingly simple explanations. For example, Darwin's theory of natural selection has struck many as elegantly simple, even though it took centuries for anyone even to frame the theory, much less to accept it as an obvious truth. Given this history, we must be modest about the credence we place in our beliefs about how complex an explanation of some phenomenon must be.

A second worry concerns how simplicity itself functions as a guide to theory choice. Simpler explanations, it is typically thought, are to be preferred other things being equal. However, other things are not always equal. It should go without saying that if a simpler explanation is flawed on some other grounds, this may be enough for us to reject it. But the argument presently under **(p.49)** consideration does not contemplate (yet) such further flaws. It instead asks us to think along the following lines: This explanation is perfectly fine so far as it goes, but I should reject a theory that yields this explanation just because the explanation is too simple. This does seem perverse. Without a more specific account of the shortcomings of a given explanation, what justifies me in thinking it is too simple?

Note that Williams's broader argument (if interpreted along these lines) is arguably immune to this criticism. This is because Williams does not stop with a comment on utilitarianism's supposedly too simple explanation; he goes on to detail what he takes to be a specific shortcoming of the utilitarian explanation. It cannot, Williams maintains, properly account for the relevance of the ground projects that inform an agent's point of view. If Williams was led to see the relevance of the agent's ground projects in part by an intuition that the utilitarian explanation was too simple, then those sympathetic to his argument should treat this as a lucky hunch. But the luck of his hunch would not leave us

justified in relying on similar hunches.

Our conclusion, then, is that a bare appeal to simplicity of explanation is unlikely to vindicate this line of argument.

IV

A more promising suggestion, and one anticipated by the foregoing discussion, is that a moral theory can become suspect for failing to account for the specific considerations that make some moral problem a hard one. A hard case is a case that involves multiple and conflicting considerations that must be reconciled or balanced, where this balancing is itself difficult to conduct. One way that a theory could end up treating (what is thought to be) a hard case as an easy case is that it might not countenance some of what are (thought to be) reasons as reasons at all. Letting R_f and R_o each stand for a reason, respectively, in favour of and opposed to A, here is such an argument form:

- (i) That A is right depends upon how R_f and R_o are balanced against one another.
- (ii) If T is true R_o is no reason at all.
- (iii) If T is true, then the fact that A is right does not depend upon how R_f and R_o are balanced against one another.

So,

- (iv) T is not true.

(p.50) One immediate and striking fact about this reading of the argument is that it is not at all clear how it might be generalized beyond the practical realm. For it is not clear what the analogue of ‘reasons for action’ would be in the context of an argument against a metaphysical theory, say. Yet (though our examples are all in the moral realm), we expect that arguments ‘from obviousness’ have been deployed in non-practical contexts.

Notice that on this argument form, ‘balancing’ takes place only once we have settled which considerations are reasons. Balancing does not include the step of deciding which considerations are reasons. This is why, if a theory implies that some consideration is not a reason, then it also implies that no balancing of such a putative reason can be called for, which is just what (iii) says.

This feature of the argument leaves an important line of reply available to the theorist accused of oversimplification. Instead of accepting (iv), he may urge, we should reject (i). We should think that some consideration that we thought was a reason—and that therefore we thought made the case a hard one—was in fact no reason at all. This reply may be bolstered by appeal to the doctrine of holism, championed by particularists and accepted by some generalists. According to that doctrine, a consideration can be a reason in one context while being no reason at all in another context. The theorist who accepts holism, therefore, need not deny that R_o is *sometimes* a reason; he need only deny that it is a reason *in this case*. And the possibility that R_o is a reason in some other cases can go some way to explaining why we might have thought it a reason in this case.

The plausibility of this reply, however, depends upon what we should say about two further issues. First, the theorist must allow that R_0 is at least sometimes a reason. Williams, for example, seems to think that the utilitarian must deny that facts about an agent's ground project are ever, themselves, reasons; any reasons would need to be located in the effects for happiness of acting contrary to one's ground projects. Second, the theorist must sow some doubts about our ability to distinguish genuine from defeated reasons.

Admittedly, the present argument is framed in a way that suggests a tight connection between reasons and rightness, that the right action is the one favoured by the balance of reasons. This connection, of course, is contested, and standard normative theories can be formulated in ways that do not implicate such a connection. For example, a standard way of defining utilitarianism is as the conjunction of a consequentialist theory of right action and a hedonistic (or, at least, welfarist) theory of value. One way of spelling this out would be as follows:

- *Consequentialism*: An action is right if and only if it promotes at least as much value as all the other actions available to the agent.
- **(p.51)** *Hedonism*: Pleasure and only pleasure is good for its own sake, pain and only pain is bad for its own sake, and the value promoted by an action is a direct function of the pleasure and pain that it promotes.
- *Utilitarianism*: Consequentialism plus Hedonism.

Crucially these doctrines *simply do not mention the concept of a reason for action*. So on this way of defining the view, there is at least no logical guarantee that the right act according to utilitarianism will also be favoured by the balance of reasons. Moreover, this strikes us as not merely being logical space which is available in the abstract but somehow implausible or strained. For one interesting view of what it is for a fact F to be a reason in favour of a given action A just is for F to explain why A promotes something of value. On this way of defining reasons for action, all sorts of facts which are not themselves facts about pleasure or pain can count as reasons for action in the right context. In particular, in the right context, the fact that an action promotes the successful pursuit of the agent's ground project can be a reason—so long as its promoting that ground project explains why the action promotes pleasure or the absence of pain.

A utilitarian who takes this route has denied (i), and with some plausibility. But this move is best seen as a sidestep, and the argument can be reformulated in a way that will require a more substantive response. This is because the utilitarian we are here imagining is still committed to the view that the right action is the action that is favoured by the balancing of *values*. And this leaves available a reframing of the argument in value, rather than reason, terms.

The defender of a normative theory may still resist by claiming that some consideration that we thought was a value—and that therefore we thought made the case a hard one—was in fact no value at all. As before, this reply may be bolstered by an appeal to holism,

though in this case holism about value, not reasons. According to holism about value, something that bears value in one context may not bear value in another context.

Our aim, we emphasize, is not to defend utilitarianism, but rather to examine the dialectical force of appeals to the morally obvious. So we will not examine here whether this line of reply can be deployed effectively by a *utilitarian*, even one willing to endorse holism. Leaving utilitarianism to the side, then, we wish to draw attention to the more general issue of our being mistaken about reasons (or values). The basic dialectic is this. Opponents of T charge that it fails to show that the rightness of A is a 'hard case'. This is because, critics claim, it fails to recognize all of the genuine and competing reasons (or values) that must be balanced both for and against A. Defenders of T urge that we not be too quickly impressed by this argument because we may have been mistaken about what is a reason (or value).

(p.52) One should doubt such disputes admit of universal solution. To adjudicate such disputes one must have recourse to the actual claims that are being advanced as reasons (or values). In principle, we see no reason that the anti-theoretical side might not prevail. So we accept that arguments from non-obviousness can sometimes succeed. We do, however, think such arguments are typically (and perhaps surprisingly) weak. Moreover, the weakness illustrates something of interest to debates over systematic and principled normative theorizing.

Over the past twenty-five years, philosophers who are at least sceptical of general normative theorizing have been among those who are most insistent that philosophers need to appreciate more carefully the subtlety and variability of the moral terrain. They have done this by highlighting the diverse ways in which facts can be morally relevant and the contingencies that affect how facts vary in relevance across cases. The development of holism—about both reasons and values—is but one instance of this trend. Bernard Williams's own rich discussions of the historical and cultural contingencies that affect moral thinking are another. There are two ways of making this last point, one of which presupposes holism and one of which does not. The first one, which we have already discussed at length, explicitly invokes holistic assumptions about reasons to distinguish different forms of normative relevance. Even a diehard atomist, though, can distinguish reasons for action from (a) things a virtuous/wise agent ought to heed in deliberation, and (b) evidence that an action is right or wrong. Insofar as arguments of this form can thereby be accused of having confused different forms of normative relevance, or perhaps of begging the question by assuming that the relevant feature just is a reason for action and not relevant in some other way, the defender of the theory in question is left room to manoeuvre.

At the risk of painting with too broad a brush, it nevertheless seems fair to say that emphasis on the subtlety and variability of the moral terrain is often thought to support anti-theoretical predilections. If the moral terrain is especially rich and nuanced, then capturing it in any general normative theory must be more difficult, if it is possible at all. Whatever the merits of this line of argument, in the case of arguments from non-obviousness the dialectical advantage runs in the other direction. The more subtle and

variable the moral landscape can be, the less sure we are entitled to be that we know which cases truly are the 'hard ones'. If we cannot be highly confident which cases are 'hard ones' then we cannot have strong reason to reject a theory on the basis that it miscategorizes a hard case. Two reminders are in order. First, 'hard cases' are not here characterized in epistemic terms, and for reasons we have already given, arguments from non-obviousness must not rely on such a characterization. Secondly, we have not tried to argue that arguments from non-obviousness cannot succeed, only that **(p.53)** they are characteristically weak. The reason for this can be put simply: The less obvious morality is, the less obvious it should be what is morally obvious.

References

Bibliography references:

Dancy, Jonathan. *Ethics Without Principles*. Oxford University Press (Oxford: 2004).

Feldman, Fred. *Doing the Best We Can: An Essay in Informal Deontic Logic*. D. Reidel Publishing Co. (Dordrecht: 1986).

Lenman, James. 'Utilitarianism and Obviousness', *Utilitas* 16 (3): 322–5 (2004).

Raphael, D. D. 'The Standard of Morals: The Presidential Address', *Proceedings of the Aristotelian Society*, 75: 1–12 (1974–5).

Williams, Bernard. 'Critique of Utilitarianism' in *Utilitarianism: For and Against* by Bernard Williams and J. J. C. Smart. Cambridge University Press (Cambridge: 1973).

Access brought to you by: University of Edinburgh