Is Moral Obligation Objective or Subjective?

MICHAEL J. ZIMMERMAN
University of North Carolina at Greensboro

Many philosophers hold that whether an act is overall morally obligatory is an 'objective' matter, many that it is a 'subjective' matter, and some that it is both. The idea that it is or can be both may seem to promise a helpful answer to the question 'What ought I to do when I do not know what I ought to do?' In this article, three broad views are distinguished regarding what it is that obligation essentially concerns: the maximization of actual value, the maximization of expected value, and the perceived maximization of actual value. The first and third views are rejected; the second view is then refined and defended. The unfortunate upshot is that there may be no very helpful answer to the question just mentioned. As to the question posed in the title of the article, the answer unsurprisingly depends on what 'objective' and 'subjective' are taken to mean.

I

Almost everyone acknowledges the ambiguity of 'ought', even when it is restricted to *moral* contexts.[1] One distinction is that between using 'ought' to express an *ideal* (as in 'No child ought to have to suffer') and using it to express an *obligation* (as in 'You ought to keep your promise'). Within the province of moral obligation, there is the distinction between what W. D. Ross calls *prima facie* obligation and what many call *all-things-considered* obligation but which I, for brevity, will call *overall* obligation. Within the province of overall obligation, many philosophers allege that there is yet a further ambiguity between what are often called *objective* obligation and *subjective* obligation. This allegation is the subject of this article.

It is with overall moral obligation that the *morally conscientious* person is primarily concerned. When one wonders what to do in a particular situation and asks, out of conscientiousness, 'What ought I to do?', the 'ought' expresses overall moral obligation. 'Ought' here is a contrary of 'wrong'. Conscientiousness precludes deliberately doing what one believes to be overall morally wrong.[2]

As I understand the debate between consequentialists and non-consequentialists, their disagreement concerns what it is that overall moral obligation consists in. I will not engage in this debate here. For

[1] An exception: Judith Jarvis Thomson, *Goodness and Advice* (Princeton, 2001), pp. 44 ff.

[2] This is not to say that conscientiousness requires deliberately doing only what one believes to be overall morally right, since on occasion one may find oneself forced to act while lacking any belief about the overall moral status of one's act.

© 2006 Cambridge University Press
Utilitas Vol. 18, No. 4, December 2006
doi:10.1017/S0953820806002159 Printed in the United Kingdom

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

the sake of convenience, I will adopt an approach that is typical of one kind of consequentialism. (I believe that my findings could be adapted to other kinds of consequentialism and to non-consequentialism, too, but I will not pursue this point.) On this approach, what we *ought* to do (in the sense that expresses overall moral obligation) is a function of the *value* of what we do, which is itself a function of some non-evaluative '*stuff*'. (My use of the term 'stuff' is intended to be metaphysically neutral.) For example, utilitarians (as portrayed by G. E. Moore[3]) believe that what we ought to do is perform that act with the highest instrumental value, which itself consists in bringing about the most favourable balance of pleasure over pain. Most consequentialists follow utilitarianism in saying that overall moral obligation involves *maximizing* the relevant value; others take a more relaxed stand.[4] Again, for the sake of convenience, I will adopt a maximizing approach. (I believe that my findings could be adapted to non-maximizing approaches, too, but I will not pursue this point either.)

We are now in a position to distinguish three broad views concerning what we ought, in the sense of overall moral obligation, to do. Using the term 'actual value' to refer to the value (whatever that may be) that resides in the stuff of moral obligation (whatever that may be), we may put these views as follows:

1. An agent ought to perform an act if and only if doing so would (uniquely) maximize actual value (among the range of options that the agent has);[5]
2. an agent ought to perform an act if and only if doing so would maximize expected value;
3. an agent ought to perform an act if and only if he (or she) believes that doing so would maximize actual value.

It is usually something like View 1 that is said to express the idea that obligation is 'objective', but not always; Bertrand Russell, for example, says this of something like View 2.[6] There is no such consensus concerning the idea that obligation is 'subjective'. Many philosophers say that something like View 2 captures this idea;[7] many others say

[3] G. E. Moore, *Ethics* (Oxford, 1912), chs. 1 and 2.

[4] For example, Michael Slote, 'Satisficing Consequentialism', *Proceedings of the Aristotelian Society* 58 (1984).

[5] The two phrases in parentheses are to be understood to be implicitly at work not only in the statement of this view (View 1) but also in the statement of all such views to follow (View 2, View 3, etc.).

[6] Bertrand Russell, *Philosophical Essays* (London, 1910), pp. 30–1.

[7] For example: Fred Feldman, *Doing the Best We Can* (Dordrecht, 1986), p. 46; Frank Jackson, 'A Probabilistic Approach to Moral Responsibility', *Logic, Methodology and Philosophy of Science VII*, ed. R. B. Marcus et al. (Amsterdam, 1986); Allan Gibbard, *Wise Choices, Apt Feelings* (Cambridge, Mass., 1990), p. 42; John Broome, 'The Structure

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

that something like View 3 does so;[8] a few say this of both View 2 and View 3.[9]

There is disagreement among philosophers whether each of Views 1–3 captures a legitimate sense of 'ought'. There is reason to think that it does. If 'ought' is ambiguous in other ways, why not in this way too? Indeed, acknowledging such ambiguity may seem helpful. Consider the plight of the conscientious person who has searched in vain for an answer to the question 'What ought I to do?' and yet finds himself forced to make a decision. In such a case he might say: 'Well, I have not been able to figure out what I ought to do. *Now* what ought I to do?'[10] If we do not distinguish between senses of 'ought', this is a puzzling question. The only answer can be, 'You ought to do whatever it is that you ought to do, whether you know what that is or not'. This is a singularly unhelpful response. If we do distinguish between senses of 'ought', though, another answer might be given: 'When you do not know what you ought (objectively) to do, you ought (subjectively) to do such-and-such'. Depending on the specification of 'such-and-such' (perhaps in terms of View 2 or View 3), this looks like it could be helpful.

Someone who answers the conscientious person's inquiry in the manner just noted would say, in response to the question posed in the title of this article, that overall moral obligation is both objective and subjective – or, more carefully, that there are both objective and subjective kinds of overall moral obligation. Despite the appeal of such a position, there is nonetheless reason to think that it is mistaken – to think, that is, that there is just one kind of overall moral obligation (whether objective, or subjective, or neither). Notice, first, that

of Good: Decision Theory and Ethics', *Foundations of Decision Theory*, ed. M. Bacharach and S. Hurley (Oxford, 1991); Mark Timmons, *Moral Theory: An Introduction* (Lanham, 2002), p. 126.

[8] For example: Russell, *Essays*, pp. 32 ff.; W. D. Ross, *Foundations of Ethics* (Oxford, 1939), pp. 146–7; H. A. Prichard, *Moral Obligation* (Oxford, 1949), p. 18; Richard B. Brandt, *Ethical Theory* (Englewood Cliffs, 1959), pp. 365–6.

[9] For example: A. C. Ewing, *The Definition of Good* (London, 1948), pp. 120–1; Derek Parfit, *Reasons and Persons* (Oxford, 1984), p. 25; Shelly Kagan, *Normative Ethics* (Boulder, 1998), p. 65; Frances Howard-Snyder, 'It's the Thought that Counts', *Utilitas* 17 (2005), p. 266. (I became aware of Howard-Snyder's article only after this article had been written.) Ewing and Howard-Snyder emphasize the distinction between Views 2 and 3; Parfit and Kagan do not. One could imagine a fourth view being proposed in this context: *An agent ought to perform an act if and only if he believes that doing so would maximize expected value.* At times, Elinor Mason (in 'Consequentialism and the "Ought Implies Can" Principle', *American Philosophical Quarterly* 40 (2003)) appears to propose this view (see especially p. 327). Compare also Torbjörn Tännsjö, *Hedonistic Utilitarianism* (Edinburgh, 1998), p. 34. I will not discuss this view further.

[10] I owe this way of framing the issue to Fred Feldman (in private correspondence). Compare Fred Feldman, 'Actual Utility, the Objection from Impracticality, and the Move to Expected Utility', *Philosophical Studies* 129 (2006), p. 67.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

regardless of just how the 'such-and-such' is specified, there is the threat of a regress. The answer just given to the conscientious person has this form: 'When you do not know what you ought$_1$ to do, you ought$_2$ to do such-and-such'. But that suggests another question, 'What ought I to do when I do not know what I ought$_2$ to do?', which in turn suggests another answer: 'When you do not know what you ought$_2$ to do, you ought$_3$ to do so-and-so'. But that suggests yet another question, with yet another answer, and so on and on. Even if we can find a way to halt the regress, the prospect is disturbing. Should we really be prepared to countenance such a proliferation of 'oughts'?

Second, there is the nagging feeling that, far from being helpful, such a proliferation of 'oughts' would succeed only in further confounding decision-making. Consider what Ross says at one point when discussing the matter of overall moral obligation:

> [W]hen people express different opinions about the rightness or wrongness of an act, the difference is often due to the fact that one of them is thinking of objective and the other of subjective rightness. The recognition of the difference between the two is therefore in itself important as tending to reconcile what might otherwise seem irreconcilable differences of opinion.

Yet he immediately goes on to say: 'But the question remains, which of the characteristics – objective or subjective rightness – is ethically the more important, which of the two acts is that which we ought to do'.[11] At first, this may seem an odd question. If Ross is prepared to recognize objective and subjective senses of 'right', why not also of 'ought'? But in fact there is clearly something to his question. We may pose it in the following more general form: once the various 'oughts' ('ought$_1$', 'ought$_2$', and so on) have been distinguished, which is the one that *really* counts? You may resist the question. 'They *all* count', you may say. But they cannot all count *equally*, for then the agent would suffer an embarrassment of riches. Suppose that, because of the quandary you are in, various 'oughts' are said to apply: you ought$_1$ to do A (even if you do not know this), ought$_2$ to do B (even if you do not know *this*), ought$_3$ to do C, and so on. If all these 'oughts' counted equally then, given that you could not do all of A, B, C, and so on, you would be in a moral dilemma. But, regardless of whether moral dilemmas should be said to be possible,[12] surely the mere fact that you are in a quandary should not be thought sufficient to put you in a dilemma.[13]

[11] Ross, *Foundations*, p. 147.

[12] See Michael J. Zimmerman, *The Concept of Moral Obligation* (Cambridge, 1996), ch. 7.

[13] Compare Earl Conee, 'Why Moral Dilemmas Are Impossible', *American Philosophical Quarterly* 26 (1989).

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

There is therefore strong pressure to say that only one of the 'oughts' that apply really counts (or at least that, if any of the others count at all, they do not count as much). But then that raises the question: which of these 'oughts' *is* the one that really counts? (Alternatively, which 'ought' *ought* you to act on?) One answer is: that 'ought' about which you are not ignorant. Thus, if you know what you $ought_1$ to do, then that is what you really ought to do. However, if you do not know what you $ought_1$ to do but do know what you $ought_2$ to do, then *that* is what you really ought to do. And so on. But that cannot be right, either. Ignorance cannot be its own reward in this way, for that would render conscientious inquiry otiose. Why bother trying to find out what you $ought_{n-1}$ to do? Just do what you $ought_n$ to do, and all will be well. But if such complacency cannot be countenanced, that returns us to our original problem. If you should not rest content with plumping for what you $ought_n$ to do but should instead (time permitting) aim at discovering and doing what you $ought_{n-1}$ to do, then you should not rest content with doing what you $ought_{n-1}$ to do but should instead aim at discovering and doing what you $ought_{n-2}$ to do – and so on, all the way back to what you $ought_1$ to do. But in that case all the 'oughts' other than '$ought_1$' would seem superfluous; we should stick with objective obligation and renounce any putative subjective obligation. But that then leaves us with the 'singularly unhelpful' response mentioned earlier.

We are thus faced with a puzzle. The question 'What ought I to do when I do not know what I ought to do?' is perfectly cogent; its cogency suggests a multiplicity of 'oughts'; this multiplicity is unacceptable, however, unless the various 'oughts' are prioritized; yet such prioritization renders all subordinate 'oughts' redundant. How is this puzzle to be resolved?

What follows is an attempt to answer this question and the more general question posed in the title of this article. The plan of the article is this: Section II – a discussion and rejection of View 3; Section III – a discussion and rejection of View 1; Section IV – an elaboration of View 2; Section V – a limited defense of View 2 and a resolution of the puzzle just raised; Section VI – a discussion of the question whether overall moral obligation is objective or subjective.

II

In order to check the credentials of Views 1–3, I will examine a series of cases, each a version of one originally devised by Frank Jackson.[14] Each version involves a physician, Jill, whose patient, John, suffers

[14] Frank Jackson, 'Decision-theoretic Consequentialism and the Nearest and Dearest Objection', *Ethics* 101 (1991), pp. 462–3. To my knowledge, the first appearance in print

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

from a minor but not trivial skin complaint. Jill has three drugs to choose from: A, B, and C. It is in fact the case that drug A will relieve but not completely cure John's condition, that drug B will completely cure it, and that drug C will kill him. Now consider this version:

> *Version 1:*
> Jill knows that giving John drug A will cure him partially, giving him drug B will cure him completely, giving him drug C will kill him, and giving him no drug will render him permanently incurable.

Here, I think, there would be little hesitation on anyone's part concerning what Jill ought to do. All else being equal, Jill ought to give John drug B. Given her knowledge about the efficacy of the alternative treatments, there is no need to consider any other 'oughts' but this. (Having said this, there is in fact reason to doubt this seemingly obvious verdict. Promissory Note no. 1: I will return to this point in Section V.)

Problems arise, though, when truth, evidence, and belief do not combine to provide Jill with such knowledge. Consider the following:

> *Version 2:*
> All the evidence at Jill's disposal indicates what in fact is the case, namely, that giving John drug A will cure him partially, giving him drug B will cure him completely, giving him drug C will kill him, and giving him no drug will render him permanently incurable. In keeping with the evidence, Jill believes that giving John drug A will cure him partially and that giving him no drug will render him permanently incurable. However, despite the evidence, she believes that giving him drug C will cure him completely and giving him drug B will kill him.

Suppose that, acting on the basis of her beliefs, Jill gives John drug C and thereby kills him. Clearly his death is very unfortunate, but should we say that Jill did not do what she ought to have done (when 'ought' expresses overall moral obligation)? Should we say, that is, that she did overall moral wrong in giving John drug C? Or should we say, as View 3 dictates, that she acted as she ought to have done?

Some would deny that Jill acted as she ought to have done, on the grounds that, in failing to heed the available evidence, she was negligent and, in acting negligently, she did wrong. But this is doubly mistaken. First, the failure to heed available evidence need not constitute negligence. (Promissory Note no. 2: I will provide a brief explanation of this in Section IV.) Second, negligent behavior need

of a case of this kind was in Donald Regan, *Utilitarianism and Co-operation* (Oxford, 1980), pp. 264–5.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

not constitute wrongdoing. (Promissory Note no. 3: I will provide a brief explanation of this in a moment.) Nonetheless, I think we should indeed reject both View 3 and its verdict in this case. I say this for three reasons.

First, if Jill was not negligent, then she is certainly not *to blame* for acting as she did; but that leaves unaffected the fact that what she did was *wrong*. We must take care to distinguish the blameworthiness of agents from the wrongness of actions; otherwise, excuses (of the sort that consist in being blameless despite having done wrong) would be a conceptual impossibility – which I take it they are not.[15]

Second, View 3 implies that all agents possess a certain sort of potential moral infallibility. If (as seems plausible) we always know, whenever we have a belief about what it would be best to do, what it is that we believe, then, if View 3 were true and we knew this, we would always know what we ought to do. But this makes a mockery of the conscientious person's inquiry into what he ought to do, implying that such an inquiry can be successfully carried out simply via introspection.

Third, View 3 implies that, on the assumption that he was doing what he believed to be best, Hitler did no wrong. But it is grotesque to think that such a belief could suffice to render such evil-doing permissible.

The three points just made are familiar, having been made in one way or another by a great many writers. They suffice, I believe, to show that View 3 is false; that is, they show that it is not the case that an agent ought to perform an act if and only if he believes that doing so would maximize actual value, *when (a) 'ought' is taken to express overall moral obligation and (b) doing what is overall morally obligatory is taken to be the primary concern of the morally conscientious person*. I stress both (a) and (b), because they are crucial to the rejection of View 3. I have no wish to deny that, if either (a) or (b) fails to apply, then View 3, or something like it, is acceptable.

Consider what A. C. Ewing says: 'We may believe . . . that the soldiers who fight against us in a war are acting wrongly in fighting, yet every reasonable person will admit that, as long as they really think they ought to fight, they ought "to obey their consciences" and fight'.[16] There is clearly a tension in what Ewing says, something that he himself is quick to recognize. How can it be that our enemies are acting wrongly if, in doing so, they are obeying their consciences and they ought to obey their consciences? Ewing's answer is that 'ought' is ambiguous, so

[15] See Michael J. Zimmerman, 'Another Plea for Excuses', *American Philosophical Quarterly* 41 (2004).

[16] Ewing, *Definition*, pp. 120–1.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

that the sense in which our enemies 'ought' to obey their consciences is different from the sense in which they 'ought not' to fight us.

I do not deny Ewing's thesis. I have denied that View 3 is true, *insofar as both (a) and (b) apply*, but it is consistent with this to claim that one ought, in some sense of 'ought' *not* governed by both (a) and (b), to do what one believes would maximize actual value. Ewing's remark may seem to suggest that he accepts this claim. In fact, he does not, and for good reason. What he does accept is the related but importantly different claim that one ought (in some sense not governed by both (a) and (b)) to do what one believes one ought (in the sense governed by both (a) and (b)) to do; or equivalently, that it is wrong (in the former sense) to do what one believes is wrong (in the latter sense).[17] I would in fact advise against putting matters in this way; such a double use of 'ought' (or 'wrong') courts confusion. Much better is to say that one is *blameworthy* if and only if one does what one believes is wrong (in the sense governed by both (a) and (b) – in the sense, that is, in which the primary concern of the conscientious person is to discover and avoid doing what is wrong).[18] This acknowledges and preserves the distinction, noted earlier, between the blameworthiness of agents and the wrongness of actions. (Redemption of Promissory Note no. 3: just as agents can do wrong but not be blameworthy, so too they can be blameworthy and not do wrong. Negligence – understood in one way, at least – consists in neglecting what one believes to be one's obligation, and thus[19] confers culpability; but the belief may be mistaken, in which case no wrong is done.[20])

[17] Note that, just as 'ought' and 'wrong' can be used to do double duty in this way, so too can 'expect'. In one sense, all that can be expected of someone is that he do what he *is* obligated to do. In another sense, all that can be expected is that he do what he *believes* he is obligated to do.

There is a complication: Ewing (*Definition*, pp. 132–3) reserves the term 'moral obligation' for a sense of 'ought' *not* governed by (b). There is a sense of 'ought' according to which the primary concern of the conscientious person is to discover and do what he ought to do. I have used the term 'overall moral obligation' for *this* sense of 'ought'. Ewing, however, says that one's 'moral obligation' is not to do what one ought, in this sense, to do, but rather to do what one *believes* one ought, in this sense, to do. This seems to me to constitute a serious misapplication of the term 'moral obligation', but the dispute is merely verbal.

[18] Actually, this is only roughly correct. The conditions of blameworthiness – at least when 'blameworthy' expresses moral culpability, the 'negative' side of moral responsibility – are more complex. In particular, a freedom condition must be satisfied. In addition, one can be blameworthy, even if one acts conscientiously, if one is to be blamed for one's conscience. (Compare Brandt, *Ethical Theory*, p. 363.) See Michael J. Zimmerman, *An Essay on Moral Responsibility* (Totowa, 1988), pp. 40 ff. for a general treatment of the conditions of blameworthiness.

[19] See the last note for a qualification.

[20] See Michael J. Zimmerman, 'A Plea for Accuses', *American Philosophical Quarterly* 34 (1997).

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

It may seem that I have been inconsistent in what I have said about conscientiousness. One acts conscientiously if one does what one believes one is overall morally obligated to do. I have said that, if one acts conscientiously, one is not to blame.[21] But I have also said that, if one acts conscientiously, one may nonetheless act wrongly; that is, one may nonetheless fail to do what one is in fact overall morally obligated to do. Yet I have characterized overall moral obligation itself in terms of conscientiousness; I have said that doing what is overall morally obligatory is the primary concern of the conscientious person. There is no contradiction. The conscientious person is not concerned with doing what his conscience dictates; he is concerned with doing what he ought – that is, is overall morally obligated – to do, and he may well worry whether what his conscience (his set of beliefs about what he morally ought to do) presently dictates is accurate. Thus he will not rest content with letting his conscience be his guide; for acting conscientiously does not guarantee that one will avoid doing wrong. This is something that the conscientious person recognizes; it is precisely his fallibility with respect to what his overall moral obligation is that drives his conscientious inquiries. However, when it comes to judging, not people's *actions* (in terms of moral obligation), but *people* (in terms of praise- and blameworthiness) in light of their actions, then we should look to whether they have acted as their consciences dictate. As long as the two types of judgment are kept distinct, there should be no suspicion of inconsistency.

III

The rejection of View 3 still leaves open the question whether overall moral obligation is properly accounted for by either or both of Views 1 and 2.

To answer this question, consider the following:

> *Version 3:*
> All the evidence at Jill's disposal indicates (in keeping with the facts) that giving John drug A will cure him partially and giving him no drug will render him permanently incurable, but it also indicates (in contrast to the facts) that giving him drug C will cure him completely and giving him drug B will kill him.

Suppose that, acting on the basis of the evidence, Jill gives John drug C and thereby kills him. Should we say that she acted as she ought to have done, which is what View 2 dictates, or should we deny this, as View 1 dictates?

[21] See n. 18 for a qualification.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

Many would deny it. Moore is a prime example. He would simply apply the distinction between wrongdoing and blameworthiness to this case too, just as he would to Version 2. He would say that, although Jill did wrong in giving John drug C, she is not to blame for doing so.[22] Is this not just as plausible an assessment of Version 3 as it is of Version 2?

For many years I thought that it was. I no longer think this, for Jackson's own version of his case undermines such an assessment.[23] That version may be put as follows:

> *Version 4:*
> All the evidence at Jill's disposal indicates (in keeping with the facts) that giving John drug A will cure him partially and giving him no drug will render him permanently incurable, but (despite the facts) it leaves completely open whether it is giving him drug B or giving him drug C that will cure him completely and whether it is giving him drug B or giving him drug C that will kill him.

Suppose that, acting on the basis of the evidence, Jill gives John drug A and thereby partially cures him. Should we say that she acted as she ought to have done? Moore and others are once again committed to denying this, but here their position is decidedly implausible. My argument for this claim is an *ad hominem* one. (We have all been taught that such arguments are to be avoided, but here it seems to me just the right type of argument to use.) Put Moore in Jill's place in Version 3. Surely, as a conscientious person, he would decide to act as Jill did and so give John drug C. He could later say, 'Unfortunately, it turns out that what I did was wrong. However, since I was trying to do what was best for John, and all the evidence at the time indicated that that was indeed what I was doing, I cannot be blamed for what I did'. But now put Moore in Jill's place in Version 4. Surely, *as a conscientious person*, he would once again decide to act as Jill did and so give John drug A. But he could *not* later say, 'Unfortunately, it turns out that what I did was wrong. However, since I was trying to do what was best for John, and all the evidence at the time indicated that that was indeed what I was doing, I cannot be blamed for what I did'. He could not say this precisely because he *knew* at the time that he was *not* doing what was best for John.[24] Hence Moore could not justify his action by appealing to View 1 – the view that he ought to maximize actual value – even though

[22] Compare Moore, *Ethics*, pp. 81–2.
[23] See the reference to Jackson (and Regan) in n. 14 above.
[24] This renders what Moore says at *Ethics*, p. 82 inadequate; so too with what Brandt says at *Ethical Theory*, p. 367.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

this is his official position.[25] On the contrary, since conscientiousness precludes deliberately doing what one believes to be overall morally wrong, his giving John drug A would appear to betray the fact that he actually subscribed to something like View 2, the view that he ought to maximize expected value.

That giving John drug A in Version 4 would maximize expected value seems clear. The expected value of an act is a function of the probabilities of its possible outcomes and the actual values associated with these outcomes. Suppose we stipulate that the actual value of providing a complete cure is 50, that of providing a partial cure is 40, that of killing John is –100, and that of rendering John permanently incurable is 0. Let us further stipulate in Version 4 that, in light of the evidence available to Jill, the probability that drug A will provide a partial cure is (for simplicity's sake) 1, that the probability that giving him no drug (call this option D) will render him permanently incurable is also 1, and that for each of drugs B and C the probability that it will provide a complete cure is .5 and the probability that it will kill John is also .5. (Promissory Note no. 4: I will have more to say about the relevant notion of probability in the next section.) Then the expected values of Jill's alternatives are as follows:

$$\mathrm{EV(A)} = (40 \times 1) = 40$$

$$\mathrm{EV(B)} = [(50 \times .5) + (-100 \times .5)] = -25$$

$$\mathrm{EV(C)} = [(50 \times .5) + (-100 \times .5)] = -25$$

$$\mathrm{EV(D)} = (0 \times 1) = 0$$

Clearly, then, giving John drug A would maximize expected value. Of course, there is reason to doubt that such exact assignments of values and probabilities can ever be given, but the beauty and power of Jackson's case is that we apparently need no such exact assignments to be confident that Jill's giving John drug A would fail to maximize actual value yet would maximize expected value. If it is agreed that she ought to give John drug A, then, this would seem to suffice to disprove View 1 while tending to confirm View 2.

There are of course responses that might be made on behalf of View 1. Let me now attend to some.

A

One response is to claim that, despite the appearances to which I have just appealed, the actual values of Jill's alternatives in Version 4 have

[25] Moore, *Ethics*, chs. 1–2.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

been miscalculated and her giving John drug A would actually be best after all, despite the fact that it would only lead to a partial cure.

But what *reason* can be given for this claim? One can of course always 'cook the books' (by fiddling either with the specification of alternatives or with the specification of their values) so that View 1 yields the intuitively correct verdict that Jill ought to give John drug A, but, unless this is done in a principled way, such a victory is merely Pyrrhic. One such principled way would be to assign actual disvalue to the running of risks, regardless of whether these risks are realized. One could then argue that the extra risk associated with Jill's giving John drug B renders her doing so actually worse than her giving him drug A, even though the former would be medically superior.[26]

This move is problematic, for two reasons. First, it seems that a Jackson-type case can be constructed that accommodates the (alleged) disvalue of risk.[27] Second, the response implies that, from the point of view of what counts regarding the determination of Jill's obligation, giving John drug B is vastly preferable to giving him drug C.[28] This seems wrong; since they are equally risky, these alternatives would seem to be on a par with one another from the point of view in question.

B

Another response that might be made on behalf of View 1 is this. Jackson's case (Version 4) indicates precisely why we must distinguish between different 'oughts'. As View 1 implies, Jill ought$_1$ (ought objectively) to give John drug B, since this is what would in fact maximize actual value. However, we can also say with perfect consistency that she ought$_2$ (ought subjectively) to give John drug A, since this is what

[26] Compare David Sosa, 'Consequences of Consequentialism', *Mind* 102 (1993), pp. 109–10.

[27] Recall the actual values recently assigned: a complete cure, 50; a partial cure, 40; death, –100; no cure, 0. Suppose that the (dis)value associated with the risk in giving either drug B or drug C is –15. Then, if we simply add in the value of risk-taking, the actual values of Jill's alternatives would be as follows: A, 40; B, 35; C, –115; D, 0. Even on View 1, then, Jill ought to give John drug A. But now suppose that Jill's evidence is such that it is not death but rather a less serious deterioration in John's health that she risks if she gives John either drug B or drug C. Let the actual value of such a deterioration be –50. Presumably, the actual (dis)value associated with risking such a deterioration will be less than that associated with risking death. Let us stipulate that this value is –8. Then the actual values of Jill's alternatives would be as follows: A, 40; B, 42; C, –58; D, 0. View 1 would then imply that Jill ought to give John drug B after all. But this is highly dubious. Note that the expected values of Jill's alternatives (when no actual value is assigned to the taking of risks) would be as follows: A, 40; B, 0; C, 0; D, 0. View 2 would thus imply that Jill still ought to give John drug A. Although the numbers just used, and the method of aggregation, are of course spurious, their being so in no way undermines the general point that they serve to illustrate.

[28] Compare the values of 35 and –115 assigned to B and C, respectively, in the last note.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

would maximize expected value. Views 1 and 2 are thus not rivals and can both be accepted.

This, too, is problematic. Views 1–3 are intended to account for what an agent ought to do, *when (a) 'ought' is taken to express overall moral obligation and (b) doing what is overall morally obligatory is taken to be the primary concern of the morally conscientious person*. I have argued that View 3 cannot be accepted under this dual condition, although it might be accepted if 'ought' is construed differently. I counseled against using 'ought' in such a way, since there is alternative terminology available that is less confusing. So too, my present use of Version 4 is intended as an argument against View 1, *under the condition that both (a) and (b) apply*. I concede that View 1 might be accepted if 'ought' is construed differently. But I would once again counsel against using 'ought' in such a way; we can simply say that Jill's giving John drug B would be *actually best* and leave it at that, avoiding any hint of confusion with the 'ought' of overall moral obligation (the 'ought' with which proponents of View 1 such as Moore have traditionally been concerned).

C

But, it may be retorted, surely the conscientious person is concerned with the actual values of his alternatives, and so it would be a mistake to say that the 'ought' of overall moral obligation is not to be understood in terms of such values.

This misrepresents my position. Expected value is in part a function of actual value; hence anyone concerned with the former will indeed also be concerned with the latter.[29] If, contrary to my advice, you insist on attaching an 'ought' to what is actually best among an agent's alternatives, we can say that the fact that Jill ought to give John drug A is in part a function of the fact that she ought to do what is actually best and give him drug B. But in saying this we must keep clearly in mind that the latter 'ought' does *not* express overall moral obligation.

In Version 4, Jill is (let us assume) very much concerned – quite correctly – with the actual values of a complete cure, a partial cure, and death. Despite this, she quite deliberately – and again quite correctly – chooses an alternative that she knows is not actually best.[30] That is precisely how any clear-headed conscientious agent would act under

[29] Compare Howard-Snyder, 'Thought', pp. 273 and 280.

[30] It is therefore important to distinguish clearly between View 2 and the following view: *An agent ought to perform an act if and only if it is most reasonable to expect that doing so would maximize actual value.* Jackson takes care to distinguish these views. (See Jackson, 'Decision-theoretic Consequentialism', p. 468.) Others do not. See, for example, Russell, *Essays*, pp. 31–2; Kagan, *Normative Ethics*, pp. 65–6.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

the circumstances. It would be appalling, morally speaking, if Jill were to give John drug B; taking such a risk would be *unconscionable*.[31]

D

But, it may again be retorted, consider Gladstone Gander, one of Donald Duck's associates.[32] Gladstone is invariably lucky; acting on his hunches, he always manages to choose that alternative which is actually best, despite being no less ignorant than the common person about the details of the situations in which he finds himself. Suppose that Gladstone were to find himself in Jill's situation in Version 4. Acting on his hunch in this case, he would give John drug B. Surely, if someone invariably maximizes actual value, then he always acts as he ought. If so, Gladstone not only would but ought to choose drug B rather than drug A. We who are not so lucky ought to choose drug A, since in choosing B we would run an unreasonable risk of harm to John. But it is implausible to say that the correct moral theory should have radically different implications for lucky and unlucky (that is, normal) agents. Rather, we should recognize that all agents have a *fundamental* obligation to maximize actual value, as View 1 declares. Moreover, all agents have a *derivative* obligation to act responsibly with respect to the production of actual value. We act responsibly in this regard when we do not run unreasonable risks with respect to the production of actual value. Gladstone would act responsibly if and only if he chose drug B, but a normal agent such as Jill would act responsibly if and only if she chose drug A. Hence Gladstone's derivative obligation coincides with his fundamental obligation, whereas there is no such coincidence in Jill's case. But Jill's fundamental obligation does coincide with Gladstone's: fundamentally, she ought to give John drug B.

I find the proposed distinction between fundamental and derivative obligation obscure, confounding, and unmotivated. It is obscure in that it is hard to see how one obligation can be 'derived' from another when the two conflict. It is confounding in that it raises once again the question of which 'ought' one ought to act on.[33] It is unmotivated in that there would appear to be no good reason to say that Gladstone has any obligation to give John drug B. His being invariably lucky

[31] Compare Gibbard, *Wise Choices*, p. 43, n. 7.

[32] I am indebted to Ingmar Persson for bringing this character to my attention and, more generally, for forcefully pressing the case for View 1.

[33] Note that, like 'ought', the term 'responsible' is treacherous. In one sense, one acts responsibly just in case one fulfills one's obligations (that is, one avoids wrongdoing). This is the sense at issue here. In another sense, one acts responsibly just in case one acts blamelessly (that is, one avoids blameworthiness). This is not the sense at issue here. Compare the discussion in the last section concerning the distinction between wrongdoing and blameworthiness.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

diminishes neither the degree nor the unreasonableness of the risk that he runs in giving John drug B. His giving John this drug would be just as unconscionable as Jill's doing so.

It could of course happen that, at some point, Gladstone's unbroken string of lucky successes entitles him to rely on hunches that, up to that point, he was not entitled to rely on. But it is precisely at that point that he is no longer lucky when he manages to do what is actually best, in that his choosing that act which in fact maximizes actual value would no longer constitute running an unreasonable risk.[34] And it is at that point that his obligations would part way with those of normal agents such as Jill.

But does this not give rise to a paradox? Suppose that Gladstone is faced with a series of situations of the sort described in Version 4. Right from the start, he relies on his hunch and chooses drug B. Doing so invariably turns out best. At some point, his string of successes entitles him to rely on his hunch. At that point, I have said, his obligation is indeed to choose drug B, whereas previously his obligation has been to choose drug A. But how can it be that, by consistently making a choice of the sort that he ought *not* to make, Gladstone manages to make it the case that he *ought* to make a choice of this very sort?

I find no paradox here. It frequently happens that one ought not to make a choice of some sort under one set of circumstances but ought to make a choice of that same sort under another set of circumstances. (Consider the choice to tell a lie, for example.) As the string of successes unfolds, Gladstone's circumstances change. True, his being faced with a choice between drugs A, B, and C is constant; but there is a dramatic change in his evidence with respect to the efficacy of choosing drug B.

E

In a final effort to defend View 1, one might propose the following. Suppose that Jill later comes into possession of new evidence that strongly indicates that drug B would indeed have cured John completely. Even if we understand and condone her having given John drug A, should we not nonetheless say now that she ought instead to have given him drug B? Indeed, is this not what she herself is likely ruefully to admit? If so, View 1 is reinstated.

There are two main ways in which the claim that Jill ought to have given John drug B may be understood: (i) she was obligated to give him drug B; (ii) she is obligated to have given him drug B. The second reading may be dismissed, for obligations cannot be retrospective

[34] He would remain lucky in another sense: he would be fortunate to have hunches whose promptings invariably coincide with the maximization of actual value.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

(contrary to what some philosophers maintain[35]). The first reading may seem to presuppose View 1, but in fact it does not; accepting the claim on that reading, therefore, does *not* provide an adequate defense of that view. On the contrary, one could say, in the spirit of View 2, that, by the light of her earlier evidence, in giving John drug A Jill acted as she ought to have done, and also that, by the light of her later evidence, she did not act as she ought to have done. Some authors are willing to countenance both claims. (Jackson is one.[36]) Note, however, that this position does *not* serve to reinstate View 1, since it does *not* imply that Jill's obligation is ever simply to maximize actual value. On the contrary, it declares Jill's obligation always to be a function not directly of the facts, but of the evidence available to her, concerning the actual values of her alternatives.

Although the position just outlined is perfectly consistent with the general theme underlying View 2, I see no reason to embrace it. I think it is a mistake to multiply 'oughts' in the proposed way. Certainly Jill would be justified in saying later that it would have been *better* to give John drug B rather than drug A, but there is no need also to say that she *ought* to have given John drug B. (That, to paraphrase Bernard Williams, is to provide the agent with one 'ought' too many.)[37] It is the evidence available to an agent *at the time* that determines what he ought *at that time* to do.[38] Later evidence is irrelevant.[39]

This dismissal of extra 'oughts' may seem too hasty, though. For the issue just raised occurs not only across times but also across agents. Consider a variation on Version 4 in which Jill's evidence is as stipulated, but Jack's evidence outstrips Jill's and strongly indicates that drug B will cure John completely. And suppose that Jill asks Jack whether she ought to give John drug B. What should Jack say?

If the 'should' in the question just raised expresses overall moral obligation, it could be that View 2 itself would imply that Jack should advise Jill to give John drug B; for that may be what would maximize

[35] See, for example, Hector-Neri Castañeda, 'The Paradoxes of Deontic Logic: The Simplest Solution to All of Them in One Fell Swoop', *New Studies in Deontic Logic*, ed. Risto Hilpinen (Dordrecht, 1981), p. 61; Feldman, *Doing the Best*, p. 43. This issue is briefly discussed in Zimmerman, *Concept*, p. 37.

[36] See Jackson, 'Decision-theoretic Consequentialism', p. 471. Compare Graham Oddie and Peter Menzies, 'An Objectivist's Guide to Subjective Value', *Ethics* 102 (1992), p. 521.

[37] Compare Bernard Williams, *Moral Luck* (Cambridge, 1976), p. 18.

[38] Compare Bart Gruzalski, 'Foreseeable Consequence Utilitarianism', *Australasian Journal of Philosophy* 59 (1981), pp. 168–9.

[39] As confirmation of this, consider the matter from *John's* perspective. Suppose that Jill had ignored the risk to John, tossed a coin ('Heads, B; tails, C'), and the coin had come up heads. Although John would no doubt be relieved to have been completely cured, it would surely be reasonable for him also to be very angry with Jill for having gambled so recklessly with his life. He would think, correctly, that she had seriously wronged him in doing so.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

expected value *for Jack*. But the more important question here is whether Jack's telling Jill that she ought to give John drug B would be *truthful*.

If the evidence available to Jack is also available to Jill, then Version 4 has in fact been misdescribed. But if Version 4 has not been misdescribed, then View 2 implies that Jack's telling Jill that she ought (that is, is overall morally *obligated*) to give John drug B would *not* be truthful.[40] This seems to me the correct verdict.[41] Jill *ought* not to give John drug B but to give him drug A instead. It would of course remain the case that, if Jack were to tell her that she would do *best* to give John drug B, he would be speaking both truly and justifiably.

IV

I have argued that both View 1 and View 3 are false. This of course does not imply that View 2 is true. The claim that in Version 4 Jill ought to give John drug A is implied by a number of theses other than the thesis that one ought in general to maximize expected value. Nonetheless, I propose that we now take the further step of assuming that View 2 is superior to all such rival theses and thus that we accept that what a conscientious person would strive to do, and what any person ought to do, is maximize expected value. I want now to refine this claim.

As noted earlier, the expected value of an action is a function of the probabilities of its possible outcomes and the actual values associated with these outcomes. I have already said that I will leave open precisely what sort of value is at stake, but, in redemption of Promissory Note no. 4, I must now say something to identify the sort of probability at issue.

It is common to distinguish between objective and subjective probability. The sort of probability at issue here is *neither* of these, as they are typically understood. Rather, as the foregoing remarks indicate, it is an *epistemic* kind of probability, having to do with the *evidence* that is available to the agent.[42] (It is the state of Jill's evidence in Version 4

[40] Unless Jack possesses an authority such that his very pronouncements have special evidential weight for Jill. In such a case, his telling Jill that she ought to give John drug B might suffice under the circumstances to make it the case that she ought.

[41] On p. 184 of *Rights, Restitution, and Risk* (Cambridge, Mass., 1986), Judith Jarvis Thomson wavers on this point, although that may in part be because she construes 'ought' differently.

[42] If causal determinism is false, certain events are such that there is an 'objective probability' less than 1 but greater than 0 that they will occur; facts of this sort may not square with the agent's evidence and may thus be irrelevant to the present discussion. 'Subjective probability', on the contrary, is typically understood in terms of the probability (of some sort) that someone ascribes, either explicitly or implicitly, to an outcome; ascriptions of this sort may also not square with the agent's evidence and may thus once again be irrelevant to the present discussion.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

that dictates that she ought to give John drug A.) We must make sure to distinguish between evidence that is *available* to someone and evidence of which that person in fact *avails* himself. Available evidence is evidence of which someone *can* and *ought*, in some sense, to avail himself. Whatever exactly this sense of 'ought' is, it is an epistemic one. It is *not* intended to be understood to express moral obligation; for that would introduce a circularity into View 2, according to which moral obligation is itself a function of expected value and, hence, of available evidence. (Redemption of Promissory Note no. 2: it is for this reason that failure to heed available evidence is not necessarily a *moral* failing and, thus, is not necessarily indicative of negligence.[43])

Assignments of epistemic probability may be understood as follows. If a proposition, p, is certain for someone, S (that is, if S is justified, epistemically, in having full confidence in p), then the probability of p for S is 1. If p is certain for S, then its negation, ∼p, is certainly false for S; in this case, the probability of ∼p for S is 0. If p and ∼p are counterbalanced for S (that is, S is justified in having some confidence in each of p and ∼p, but no more confidence in one than in the other), then the probability of each of p and ∼p for S is .5.[44] If S is justified in having greater confidence in p than in ∼p, then the probability of p for S is greater than .5 and the probability of ∼p for S is less than .5; in such a case, p may simply be said to be probable for S, and ∼p improbable.[45] (It may frequently be the case that S's situation is such that the probability of p for S cannot even in principle be assigned a precise number between 0 and 1, but for the sake of simplicity I will here assume otherwise.[46]) I will use the term 'evidence' to refer to whatever it is that justifies someone's being more or less confident in a proposition. (I will not inquire into the nature of evidence.) The stronger S's *total*

[43] I do not mean to deny that there is a moral question about what evidence to seek, use, and so on, or that the answer to this question will vary according to what values are at stake. My point is that the notion of available evidence is itself not a moral one but an epistemic one. The '-able' indicates not simply evidence that can be accessed but evidence that is in some sense epistemically worthy of access, so that, if one's beliefs do not comport with one's evidence, one has made an epistemic mistake, whether or not one has made a moral mistake.

[44] The stipulation that S is justified in having some confidence in each of p and ∼p ensures that the equal probabilities of these propositions are not derived simply from S's being wholly ignorant about them. Compare D. H. Mellor, *Probability* (London, 2005), pp. 27–9, on the principle of indifference.

[45] Compare Roderick M. Chisholm, *Theory of Knowledge*, 3rd edn. (Englewood Cliffs, 1989), ch. 2, for a somewhat similar account. However, Chisholm writes in terms of variations in the degree (or level) of the justification that there may be for a belief, rather than in terms of variations in the degree of belief (or confidence) that may be justified. Also, he does not put the concept of probability to use as I have done.

[46] There is, in addition, the question of how epistemic probability of the sort outlined here relates to the standard probability calculus. Certain of the standard axioms do not apply in the present context. I cannot pursue the issue further here.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

evidence for p, the greater the confidence S is justified in having in p. (Presumably, in order for S to *know* p, p must be probable for S, but more than that I will not venture to say – not only because of Gettier-type concerns having to do with defective evidence, but also because of lottery-paradox-type concerns having to do with very high probabilities that seem nonetheless insufficient for knowledge.) It can clearly happen that the evidence available to S regarding p is different from the evidence available to S* regarding p, so that the probability of p for S is different from the probability of p for S*. Just what it is for evidence to be available to someone, however, is a difficult issue that I will also leave to one side.[47]

Let us now return to Version 4. There the verdict was that Jill ought to give John drug A, because doing so would maximize expected value – any other choice would be unacceptably risky. But this verdict may have been too hasty. It may seem obviously correct, because it may seem obvious what the actual values at stake are. Yet it could be that it is *not* obvious to Jill what these values are. If so, this would make a difference to what she ought to do. Let me explain.

In the last section I stipulated, for the sake of illustration, the following actual values: providing a complete cure, 50; providing a partial cure, 40; killing John, −100; rendering John permanently incurable, 0. On this basis, the following expected values were assigned:

$$\mathrm{EV(A)} = (40 \times 1) = 40$$

$$\mathrm{EV(B)} = [(50 \times .5) + (-100 \times .5)] = -25$$

$$\mathrm{EV(C)} = [(50 \times .5) + (-100 \times .5)] = -25$$

$$\mathrm{EV(D)} = (0 \times 1) = 0$$

These calculations reflect the (unrealistic) presupposition that the evidence available to Jill is such that it is certain for her that giving John drug A will partially cure him and that giving him no drug will render him permanently incurable, while, for each of drugs B and C, she is justified in having no more confidence in the proposition that giving John the drug will cure him completely than in the proposition that it will kill him. Note that, although the calculations are sensitive to deficiencies in Jill's evidence regarding the possible outcomes of her options, they are not sensitive to deficiencies in her evidence regarding the possible values of these outcomes. It is as if all the relevant values are being treated as being certain for her, even if some of the relevant

[47] For a valuable discussion of this elusive notion, see Richard Feldman, 'Having Evidence', *Philosophical Analysis*, ed. D. F. Austin (Dordrecht, 1988).

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

outcomes are not. But this could well be false. It could happen, for instance, that all the relevant outcomes are certain while some of the relevant values are not. For example, suppose that Jill has a pill that she can dispense to either John (act E) or to Jane (act F), but not to both. The pill is certain to induce a partial recovery in John but a complete recovery in Jane. The problem is that, although John is a human being, Jane is not – she is a hamster (Jill is both a physician and a vet) – and Jill's evidence concerning the relative values of the lives of humans and hamsters is equivocal. It is certain for her that the value of a partial recovery in John is 100, but her evidence concerning the value of a complete recovery in Jane is divided. She is justified in having some confidence in the proposition that hamsters' lives are considerably less valuable than humans', and thus to set the value of a complete recovery in Jane at 20. But she is justified in having equal confidence in the proposition that such an assignment of value is merely speciesist, and that the value of a complete recovery in Jane is actually 120. Under the circumstances, what I will call the *expectable values* of Jane's alternatives may be computed as follows:

$$E^*V(E) = (100 \times 1) = 100$$

$$E^*V(F) = [(20 \times .5) + (120 \times .5)] = 70$$

(Whereas the expected value, EV, of an act is a function of the probabilities of its possible outcomes and the *actual* values associated with these outcomes, the expectable value, E^*V, of an act is a function of the probabilities of its possible outcomes and the *probable* values associated with these outcomes.) The general lesson of Version 4 is that an agent's overall moral obligation is a function, not directly of the outcomes of his options, but of the evidence available to him concerning these outcomes. Such evidence may be empirical (as in Version 4, where the relevant evaluative facts are implicitly taken as certain) or evaluative (as in the present case of John and Jane, where the relevant empirical facts are implicitly taken as certain). In his own discussion of his case, Jackson pays attention only to empirical probabilities and ignores evaluative probabilities. This is unwarrantedly one-sided. We should take both kinds of probabilities into account.[48] If we do so (and if we have already

[48] In his 'Consequentialism and the Nearest and Dearest Objection', *Minds, Worlds, and Conditionals*, ed. Ian Ravenscroft (Oxford, 2006), and 'Moore on the Right, the Good, and Uncertainty', *Metaethics after Moore*, ed. Terence Horgan and Mark Timmons (Oxford, forthcoming), Michael Smith notes that there is no good reason to claim that overall moral obligation is a function of empirical but not of evaluative probabilities. However, he concludes that neither kind of probability is relevant, rather than that both kinds are. In doing so, he appears not to recognize that the Moore-type move of (correctly)

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

taken the 'further step' mentioned at the outset of this section), then we will arrive at the view that overall moral obligation is a matter of what I have called expectable value.[49] Thus I now propose that, instead of View 2 as stated, we consider moving to the following version of the general idea that it represents:

2a. An agent ought to perform an act if and only if doing so would maximize expectable value.

It seems clear, however, that this cannot be the final move, since the evidence pertaining to the expectable values of one's alternatives could be just as defective as the evidence pertaining to their actual values. Version 4 (in conjunction with the 'further step') shows us that, when one's evidence is such that maximizing actual value conflicts with maximizing expectable value, we should sacrifice the former in favour of the latter. But surely a similar case can be constructed in which one's evidence is such that maximizing expectable value conflicts with maximizing expectable expectable value, in which case we should once again sacrifice the former in favour of the latter.

To see the problem, consider the original assignment of values to Version 4:

	Actual Value	Expected Value
A	40	40
B	50	−25
C	−100	−25
D	0	0

Let us assume, unrealistically but for the sake of simplicity, that it is certain for Jill that the value of providing a complete cure is 50, that of providing a partial cure is 40, that of killing John is −100, and that of rendering him permanently incurable is 0. Then the expectable values will match the expected values. That is:

distinguishing between wrongdoing and blameworthiness is not an adequate response to Version 4.

[49] The usual calculation for expected value goes as follows: for each possible outcome, O_i, of an act, A, multiply its probability, given A, by its (actual) value, and then sum these products. The resulting formula for the expected value of A is this: $\sum_i \text{prob}(O_i/A) \times V(O_i)$. The calculation for expectable value introduces a further variable, in that each possible outcome has a number of possible values, V_j, whose probabilities must be accommodated. The resulting formula for the expectable value of A is this: $\sum_i \text{prob}(O_i/A) \times \sum_j \text{prob}(O_i \text{ has } V_j) \times V_j$. (For this formula to be applicable, it must be assumed that, if V_j is infinite, then $\text{prob}(O_i \text{ has } V_j) = 0$.)

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

	Actual Value	Expectable Value
A	40	40
B	50	−25
C	−100	−25
D	0	0

Let us further unrealistically assume that not only is it the case that the values of providing a complete cure, and so on, are certain for Jill, but also it is certain for her that they are certain for her. But now let us suppose that the relevant empirical probabilities are *not* certain for her. For example, I have stipulated (unrealistically) that it is certain for Jill that giving John drug A will partially cure him; that is, that

$$\text{prob}(\text{partial cure}/\text{A}) = 1.$$

Suppose, however, that it is not the case that

$$\text{prob}[\text{prob}(\text{partial cure}/\text{A}) = 1] = 1.$$

Suppose instead that

$$\text{prob}[\text{prob}(\text{partial cure}/\text{A}) = .8] = .5, \text{ and}$$
$$\text{prob}[\text{prob}(\text{partial cure}/\text{A}) = .4] = .5.$$

We could then say that

$$\text{prob}_2(\text{partial cure}/\text{A}) = [(.8 \times .5) + (.4 \times .5)] = .6.$$

Suppose further that

$$\text{prob}_2(\text{death}/\text{A}) = .4.$$

Then (given that there are no evaluative uncertainties and only empirical uncertainties), the expectable expectable value of Jill's giving John drug A is this: $[(40 \times .6) + (-100 \times .4)] = -16$. Suppose that similar calculations resulted in the following values for B, C, and D, respectively: −30, 10, and 5. The situation would then be this:

	Actual Value	Expectable Value	Expectable Expectable Value
A	40	40	−16
B	50	−25	−30
C	−100	−25	10
D	0	0	5

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

And here is the point. View 1 says that Jill ought to maximize actual value and so implies that she ought to give John drug B. I have argued that this is mistaken. View 2a says that Jill ought to maximize expectable value and so implies that she ought to give John drug A. But surely this is mistaken too. Ought she not rather to maximize expectable expectable value and so give John drug C?

But of course that cannot be the final move, either. What if maximizing expectable expectable value conflicts with maximizing expectable expectable expectable value? Clearly a regress threatens, and this is troubling. There is a solution, however. Let us call what I have so far called expectable value 'expectable value at level 1'. Let us call what I have called expectable expectable value 'expectable value at level 2'. And so on. Now, no human agent is such that there can ever be an infinite number of levels of evidence pertaining to expectable value; on the contrary, on any occasion the number of such levels is likely to be very small, because one's being justified in having any degree of confidence in a proposition requires that one grasp that proposition, and propositions involving more than just a few levels of such evidence are likely to be beyond anyone's grasp. Thus for every human agent there will always be a level of evidence, L, such that maximization of expectable value at that level does not on that occasion conflict with maximization of expectable value at any higher level, either because, owing perhaps to the agent's cognitive limitations, there is no such higher level or because, if there is such a level, any such level is one at which what maximizes expectable value at it also maximizes expectable value at level L. Let us call L the agent's 'definitive' level of evidence on that occasion. I then propose that, instead of View 2a, we accept the following (still rough) version of the general idea it represents:

2b. An agent ought to perform an act if and only if doing so would maximize expectable value at his definitive level of evidence.[50]

Even though still rough, View 2b is unfortunately rather complicated.[51] I'm afraid, though, that that just is the lesson of Jackson's case. Overall moral obligation is not as simple a matter as Moore would have it or as we might wish it to be.

[50] If, for some non-human agent in some situation, there is no definitive level of evidence, then that agent has no moral obligation in that situation.

[51] Calculation of expectable value at levels greater than 1 is a complex matter, especially since the relevant empirical and evaluative probabilities can vary independently of one another. The fact that evidence rarely if ever yields precise probabilities (see the discussion above of assignments of epistemic probability) simplifies the calculation in one way, since one need not worry about computing exact products and sums. But in another way it makes the calculation more complex, since in principle it involves the difficulty of making fuzzy measurements (by means, perhaps, of appealing to intervals rather than to determinate numbers).

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

V

I have all along said that Views 1–3 have been stated roughly; I have said this even when View 2 is put in the terms of View 2b. The reason for my saying this can now be given. To say that an act would uniquely maximize a certain value (whether actual, expected, or expectable) among the range of options that the agent has is to say that it is superior in respect of such value to its alternatives in that range. There are, however, serious problems involved in trying to spell out the relevant sense of 'alternative'.[52] These problems can, I believe, be satisfactorily resolved by adopting what has come to be called a 'possibilist' account of obligation.[53] A more precise statement of what I take to be the correct criterion of overall moral obligation would therefore involve a modification to View 2b that accommodates possibilism. There is no need (and no room) to undertake such a modification here, however.[54] Instead, let me attend to three objections.

A

The first is this. I myself objected earlier to View 3 that it had the 'grotesque' implication that Hitler did no wrong (as long as he was doing what he believed to be best). A similar objection may be raised against View 2b. Suppose that one is faced with the choice between causing great pain to some innocent people or refraining from doing so, and suppose that it would actually be best to refrain from doing so, but that for some reason (because one's empirical or evaluative evidence is misleading) it would maximize expectable value (at one's definitive level of evidence) to cause the pain. Is it not 'grotesque' to think that under such circumstances one's overall moral obligation is to cause the pain?

Unsurprisingly, I deny the charge. The fact is, any plausible moral view will have implications that may initially be unpalatable. For example, a proponent of View 1 should recognize, as Moore himself does, the distinction between wrongdoing and blameworthiness. In acknowledging the possibility of excuses, one is acknowledging the possibility of *blameless wrongdoing*. Hitler committed enormous evil. An application of View 1 will likely yield the conclusion that he therefore did great wrong.[55] But even on View 1 it could be that Hitler was blameless. This perhaps unwelcome possibility should not be dismissed

[52] Compare Hector-Neri Castañeda, 'A Problem for Utilitarianism', *Analysis* 28 (1968); Lars Bergström, 'Utilitarianism and Alternative Actions', *Noûs* 5 (1971).

[53] See Zimmerman, *Concept*, ch. 2.

[54] I address the issue in Michael J. Zimmerman, 'The Relevance of Risk to Wrongdoing', *The Good, the Right, Life and Death*, ed. Richard Feldman, Kris McDaniel, Jason Raibley, and Michael J. Zimmerman (Aldershot, 2006).

[55] Whether it does yield this conclusion depends on just how the view is developed.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

lightly. (On my view, whether Hitler was to blame depends in part on whether he non-culpably believed that he was doing no wrong. I see no reason to believe that he did.) An application of View 2b might yield the conclusion that Hitler did no wrong. (Whether he did so would depend in part on the evidence available to him concerning the relative values of sparing and taking the lives of millions of innocents. I see no reason to believe that his evidence differed significantly from ours in this regard.) View 2b thus opens up the perhaps unwelcome possibility of *wrongless evil-doing* in a way that View 1 does not. This of course does nothing to diminish the evil in question and should not be dismissed as grotesque, any more than the possibility of blameless wrongdoing should be dismissed as grotesque. Notice that the allegation that View 3 is grotesque was based on quite different grounds, namely, that it implies that whether one does wrong is a function of what one *does* believe (as opposed to what one is *justified* in believing).

B

The second objection is this. In arguing for View 2b to the exclusion of other views such as Views 1 and 3, I have done nothing to accommodate the propriety of the question, raised in Section I, 'What ought I to do when I do not know what I ought to do?' This question presupposes two senses of 'ought', but View 2b provides an account of only one sense.

In an attempt to resolve this issue, let me first draw your attention to the 'opposite' question, namely, 'What ought I to do when I *do* know what I ought to do?' Everyone will agree that the answer to this question is, 'You ought to do that which (you know) you ought to do'. Notice that the following question is different: 'What ought I to do when I know what will maximize actual value?' Here we cannot assume that everyone will agree what the answer is. We know what answer it is that View 1 implies: 'You ought to do that which (you know) will maximize actual value'. But what answer does View 2b imply?

Return to Version 1, which goes as follows:

> *Version 1*:
> Jill knows that giving John drug A will cure him partially, giving him drug B will cure him completely, giving him drug C will kill him, and giving him no drug will render him permanently incurable.

I said earlier that there would be little hesitation on anyone's part in saying that, in this version of the case, Jill ought to give John drug B, but I also said that there is nonetheless reason to doubt this verdict. In redemption of Promissory Note no. 1, this reason can now be given. Even if Jill knows that giving John drug B will cure him completely and that giving him drug C will kill him, that by itself does not imply that she ought to give him drug B rather than drug C – not if, for

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

example, her available evidence indicates that killing him is more valuable than curing him; under those circumstances, View 2b implies that Jill ought to give him drug C rather than drug B. But presumably these circumstances are ones under which Jill does not *know* what will maximize actual value precisely because her available evidence concerning the relevant values is defective. What, though, if her evidence is not defective in this respect? Could there still be reason to deny that she ought to give John drug B?

Suppose that Jill's evidence concerning the relevant values is accurate and, moreover, that she is justified in having a high degree of confidence in the proposition that giving John drug B will cure him completely. But suppose that she also has some reason to suspect that giving him drug B will kill him, and thus that she cannot be certain that doing so will cure him completely. Under such circumstances, giving John drug B might be unacceptably risky, so that what Jill ought to do is give him drug A instead. Now consider this question: is it compatible with the circumstances just outlined that Jill *knows* that giving John drug B will cure him completely? (That is, could the high probability of the proposition suffice under the circumstances for her knowing it, despite the fact that she also has evidence against it?) If the answer to this question is 'Yes', then View 2b does not in general imply, as View 1 does, that the answer to the question 'What ought I to do when I know what will maximize actual value?' is 'You ought to maximize actual value'. I suspect, however, that we should *not* say that, under the circumstances, Jill knows that giving John drug B will maximize actual value. (I am in no position to insist on this, though, having declined to specify just what the relation is between epistemic justification and knowledge.) If this is right, then the answer that View 2b in general implies to the question 'What ought I to do when I know what will maximize actual value?' would seem to be the same as that implied by View 1: 'You ought to maximize actual value'. This is because, under the circumstances envisaged, maximization of actual value and maximization of expectable value apparently coincide.

Let us now turn to the question that *is* raised by the puzzle: 'What ought I to do when I do *not* know what I ought to do?' To begin with, we should distinguish this question from the following: 'What ought I to do when I do not know what would maximize actual value?' Notice that the latter question does not even hint at a multiplicity of 'oughts'. Notice also that, while the answer that View 1 provides ('You ought to do that which would maximize actual value') is not at all helpful, the answer that View 2b provides ('You ought to do that which, at your definitive level of evidence, would maximize expectable value') could be more helpful. It is plausible to contend that it is often (and perhaps always) very difficult (and perhaps impossible) to determine what would maximize

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

actual value. (Whether this is indeed the case depends on what kind of value is at stake and on what kind of stuff this value supervenes on. As I said earlier, I will not investigate this issue.) But even if this is so, determining what would maximize expectable value could well be easier, since this only involves attending to the evidence that is available.

I say 'only', but this must not be misconstrued. It could well be that on many occasions attending to the available evidence remains a difficult task. It should not be thought that the criterion of moral obligation that View 2b supplies constitutes a readily applicable decision procedure.[56] There *is* a procedure that it provides ('Peruse the available evidence regarding your alternatives and their possible outcomes; assess the expectable values of these alternatives; then do that which, at your definitive level of evidence, has the highest expectable value'), but how easy it is to apply this procedure is another matter. Indeed, to apply it with precision will surely always be very difficult,[57] but precision is not always necessary in order for one to be confident that one has applied the procedure successfully. (Consider Jackson's own case, that is, Version 4. There it is easy to figure out what would maximize expectable value, even if the relevant probabilities cannot be precisely determined.[58] Hence the case's beauty and power. Still, it must also be admitted that the case is an artificially simple one.)

Suppose that someone has tried and failed to apply the decision procedure provided by View 2b and that the time to act has come. He may say, 'I accept View 2b and have tried to discover what, according to that view, I ought to do under the present circumstances. But I have not managed to figure this out. *Now* what ought I to do?' Insofar as this question is asked out of conscientiousness, we might construe it as a request for help with achieving a sort of prospective 'damage control'. So understood, the question is asked out of a fear of doing wrong and

[56] Concerning the relation between criterion and decision procedure, see: Eugene R. Bales, 'Act Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?', *American Philosophical Quarterly* 8 (1971); R. M. Hare, *Moral Thinking* (Oxford, 1981); Alastair Norcross, 'Consequentialism and the Unforeseeable Future', *Analysis* 50 (1990); Robert L. Frazier, 'Act Utilitarianism and Decision Procedures', *Utilitas* 6 (1994); James Lenman, 'Consequentialism and Cluelessness', *Philosophy and Public Affairs* 29 (2000); Timothy Chappell, 'Option Ranges', *Journal of Applied Philosophy* 18 (2001); Jonas Gren, *Applying Utilitarianism: The Problem of Practical Action-Guidance* (Gothenburg, 2004); Feldman, 'Actual Utility'.

[57] This is appropriately stressed in Feldman, 'Actual Utility'. Compare also Mark Strasser, 'Actual versus Probable Utilitarianism', *Southern Journal of Philosophy* 27 (1989); Gren, *Applying Utilitarianism*. Compare also n. 51 above. Contrast Ewing, *Definition*, who declares (on p. 128) that 'in the light of probabilities we can determine with relative ease and even sometimes with fair certainty what an agent ought to do'.

[58] Compare a case given in Howard-Snyder, 'Thought', p. 276.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

in the hope that any wrongdoing can be minimized. What help can we give?

In an effort to help, we might try invoking the distinction between unconditional and conditional obligation. For an illustration of this distinction, suppose that it is Matt's overall moral obligation (as determined by the account given in View 2b) to attend a meeting on the first floor of his office building.[59] This is to be understood as an unconditional obligation. Even though he has this obligation, it may nonetheless be the case that Matt also has an overall moral obligation (as determined by an extension of the account given in View 2b) to attend a meeting on the second floor (due to take place at the same time as the meeting on the first floor), *if* he fails to attend the meeting on the first floor. This would be a conditional obligation. It is commonplace to observe that a conditional obligation does not in general warrant the detachment of an unconditional obligation simply in virtue of its condition's being satisfied.[60] If it did, then, if Matt were in fact to be about to fail to attend the meeting on the first floor, he would have an unconditional obligation to attend the meeting on the second floor. But this obligation would conflict with his obligation to attend the meeting on the first floor, and hence Matt would be in a moral dilemma. But surely there is no dilemma. Matt's only unconditional obligation is to attend the meeting on the first floor. If he does that, then, even though he must of course thereby fail to attend the meeting on the second floor, he does no wrong whatsoever.

The commonplace observation just mentioned is simplistic. We can and should detach unconditional obligations from conditional obligations, while avoiding dilemmas, by invoking levels of obligation. Let us say that Matt ought$_1$, that is, has a *primary* obligation, to attend the meeting on the first floor; that he ought$_2$ to attend the meeting on the second floor, if he fails to attend the meeting on the first floor; that he ought$_3$ to attend a meeting on the third floor, if he fails to attend either of the first two meetings; and so on. We can say that, if Matt in fact fails to attend the meeting on the first floor, then he has a *secondary un*conditional obligation (that is, he *un*conditionally ought$_2$) to attend the meeting on the second floor; and so on.[61] Suppose that Matt does not attend the meeting on the first floor. Then he does (unconditional)

[59] This example and the ensuing discussion derive from Zimmerman, *Concept*, ch. 4.

[60] Such detachment was first given the useful label 'factual detachment' in Patricia S. Greenspan, 'Conditional Oughts and Hypothetical Imperatives', *Journal of Philosophy* 72 (1975).

[61] Factual detachment may in general be represented as follows: $O(q/p)\&p \vdash O(q)$. My claim is this: such detachment is valid when the obligations at issue are on the same level. Thus, whereas we should not endorse $O_2(q/p)\&p \vdash O_1(q)$, we should endorse, for any n, $O_n(q/p)\&p \vdash O_n(q)$. For further discussion, see Zimmerman, *Concept*, ch. 4.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

wrong. If he fails also to attend the meeting on the second floor, he *compounds* this (unconditional) wrongdoing. If he fails also to attend the meeting on the third floor, he compounds it still further; and so on. Notice, however, that the relevant unconditional subsidiary (that is, non-primary) obligations are detached if but *only* if the relevant conditions are satisfied. If Matt attends the meeting on the first floor, he does no unconditional wrong whatsoever. As it stands, View 2b implicitly concerns only primary unconditional obligation. It can be extended to cover subsidiary obligation. In that way levels of 'ought' can easily be accommodated, and moreover there is no threat of a vicious regress; for, like the number of floors in Matt's building, the number of levels of obligation will be finite.

It may seem that we could apply this approach to the present problem. We might say: if you do not know what your primary obligation is, then you have a secondary obligation to do such-and-such; if you do not know either what your primary obligation is or what your secondary obligation is, then you have a tertiary obligation to do so-and-so; and so on. But, although View 2b can and should indeed be extended to accommodate subsidiary obligation, doing so will unfortunately *not* help us solve the present problem. This is because ignorance of one's obligation, at whatever level, does *not* generate some further subsidiary obligation. Consider Matt. His primary obligation is to attend the meeting on the first floor; doing so is, we may say, 'deontically supreme', in that it is superior (from the point of view of what determines obligation) to any of his other options. To say that he has a secondary obligation to attend the meeting on the second floor, if he fails to attend the meeting on the first floor, is to say that, if the option of attending the meeting on the first floor is abstracted from his range of options and thus discounted, then the option of attending the meeting on the second floor becomes deontically supreme. In order for this conditional obligation to trigger a corresponding unconditional obligation, it must be the case that the option that is abstracted does *not* occur. But then this approach cannot be applied to our present problem, precisely because ignorance of an action's being obligatory does not entail non-performance of that action. Matt may be ignorant of his obligation to attend the meeting on the first floor, but *that* does not trigger an obligation to attend the meeting on the second floor. Despite his ignorance, Matt might nonetheless attend the meeting on the first floor. If he does, then he does no wrong whatsoever. It is therefore *not* in general the case that, if you do not know what your primary obligation is, then you have a secondary obligation to do something else.[62]

[62] More generally still, it is not the case that $O_n(p) \,\&\, \sim K[O_n(p)] \vdash O_{n+1}(\sim p)$. Again, the reason is that, even if you do not know that you ought$_n$ so to act that p (i.e. $\sim K[O_n(p)]$), this

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

Though initially tempting, invoking subsidiary obligation in the effort to resolve the puzzle posed in Section I thus strikes me as futile. And I can find no better way to handle the request for 'damage control'. I am afraid that, if someone has tried and failed to discover what act View 2b requires of him, and 'crunch time' has come and he asks you what he ought now to do, then if you also do not know what would, at his definitive level of evidence, maximize expectable value for him, all you can say is this: 'You ought to do that which, at your definitive level of evidence, would maximize expectable value, whatever that may happen to be'. There is no other 'ought' to invoke, and at least you will have spoken truthfully, if not helpfully.

Or rather: there is no other 'ought' *that expresses overall moral obligation* to invoke. You can of course always resort to the familiar answer 'You ought to do whatever you think you ought to do' (or perhaps the somewhat less familiar answer 'You ought to do whatever you think gives you the best chance of doing what you ought to do'), but in general this will be true only if the first 'ought' does *not* express overall moral obligation (given that the second 'ought' does).[63] The answer can then be understood as equivalent to 'You will avoid blameworthiness if you do whatever you think you ought to do (or, perhaps, if you do whatever you think gives you the best chance of doing what you ought to do)', in which case, as noted in Section II, I think it can be accepted (with qualifications[64]). But, even if acceptable, this answer will unfortunately be of little help. Even if the person asking the question does already believe, of some act A, that he is obligated to do A (or that doing A will give him the best chance of fulfilling his obligation) – which he may well not, since he is seeking your advice on the matter – telling him

is consistent with your nonetheless so acting that p and thus with your doing what you ought_n to do; and if you do do what you ought_n to do, then no unconditional obligation_{n+1} to act otherwise will arise.

Here is an alternative route to the same conclusion. The following principle is plausible: $O_n(p) \,\&\, O_{n+1}(\sim p/q)] \vdash O_n(\sim q)$. (For an account of moral obligation that sanctions this principle, see Zimmerman, *Concept*, ch. 4.) Now suppose, for purposes of *reductio ad absurdum*, that (1) you ought_1 so to act that p and that (2) you ought_2 so to act that $\sim$p, given that you do not know that you ought_1 so to act that p. In light of the principle just mentioned, it would follow that (3) you ought_1 so to act that you know that you ought_1 so to act that p. But surely (3) is, as a general rule, false: we do not in general have an obligation to know what our obligations are; that is, we are sometimes justifiably ignorant of our obligations. Thus, if (1) is true, (2) is false; and if (2) is false, no secondary obligation so to act that $\sim$p will arise from your ignorance that you ought_1 so to act that p.

[63] The 'oughts' could both express overall moral obligation and your answer could still be true, since it can of course happen that a person is both obligated to do A and believes that he is obligated to do A. But your speaking truthfully would then be sheer happenstance, since *ex hypothesi* both you and he are ignorant of the fact that his belief about what he ought to do is accurate.

[64] See n. 18 above.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

that he will avoid *blameworthiness* if he acts on this belief may provide him with some sort of consolation, but it simply does not address his concern with avoiding *wrongdoing*.

Here, then, is how I would respond to the puzzle posed in Section I. The question 'What ought I to do when I do not know what I ought to do?' is indeed perfectly cogent, but its cogency does not entail a multiplicity of 'oughts' and we should resist any suggestion that there is such a multiplicity (barring the 'oughts' of subsidiary obligation – but they are not pertinent here).[65] The best answer that we can give to it, helpful or not, is simply this: 'You ought to do whatever it is that would, at your definitive level of evidence, maximize expectable value'. If the person who asks the question believes that there is a better answer, he is deluding himself.

C

This response to the puzzle is perhaps more a dissolution than a solution of it and may for that reason be disappointing.[66] It may provoke the following third and final objection: 'If the unhelpfulness of the response "You ought to do that which, at your definitive level of evidence, would maximize expectable value" (made to someone who has not been able to figure out what would maximize such value) is not sufficient reason to reject View 2b, then the unhelpfulness of the response "You ought to do that which would maximize actual value" (made to someone who has not been able to figure out what would maximize such value) is not sufficient reason to reject View 1. But in that case there is no need to move from View 1 to View 2b in the first place'. This, however, misrepresents the reason for moving from View 1 to View 2b. That reason was *not* to find a helpful response to the question 'What ought I to do when I do not know what would maximize actual value?'[67] Rather, the move was dictated by the recognition that Jackson's case (Version 4) shows quite clearly that it is not in general the case that one ought, in the sense that expresses overall moral obligation, to

[65] I think we should also resist any suggestion that there is a multiplicity of 'rights' and 'wrongs'. (Contrast Ross, quoted in Section I above.) One acts wrongly just in case one fails to do what one ought; one acts rightly just in case one avoids acting wrongly. Just as there is only one 'ought' in the sense of overall (unconditional, primary) moral obligation, so too there is only one related sense of 'right' and only one related sense of 'wrong'. Compare Ewing, *Definition*, p. 118.

[66] On pp. 237 ff. of 'Prichard on Duty and Ignorance of Fact', *Ethical Intuitionism*, ed. Philip Stratton-Lake (Oxford, 2002), Jonathan Dancy proposes a solution to our problem that is certainly not a dissolution of it but which, at the same time, correctly (as I have argued) avoids multiplying senses of 'ought'. I believe the solution fails, for reasons that there is no space to discuss here. (Briefly, though: his solution appears both to violate the principle that 'ought' implies 'can' and to entail that moral dilemmas arise in situations in which it is implausible to think that they do.)

[67] Contrary to the suggestion made in Feldman, 'Actual Utility', pp. 50 and 62–3.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

maximize actual value. If the response 'You ought to do that which, at your definitive level of evidence, would maximize expectable value' is helpful to someone who has not been able to discover what would maximize actual value, that is simply a bonus, and not the purpose, of the move from View 1 to View 2b. That the latter view does not itself furnish a helpful response to the further question 'What ought I to do when I do not know what would maximize expectable value at my definitive level of evidence?' gives no reason to reject it in turn.

VI

What, finally, should be said, in the light of View 2b, in response to the question 'Is moral obligation objective or subjective?' The answer of course depends on what 'objective' and 'subjective' are taken to mean. In one common sense, obligation may be said to be objective if it is possible for someone to be mistaken about what he ought to do, and subjective otherwise. On this understanding, View 2b implies that obligation is objective. (As is not the case with View 3, even if one knew that View 2b was true, one could easily fail to know what one ought to do.) This seems to me exactly what we should say. Conscientiousness presupposes that obligation is objective, in this sense. If obligation were not objective, conscientiousness would thus be *fundamentally* delusional, which it surely is not.

In another common sense, obligation may be said to be objective if it is not even in part a function of the agent's mental state, and subjective otherwise. On this understanding, View 2b implies that obligation is subjective. (Note that, even though View 2b implies that what one ought to do is in part a function of what one is *justified* in believing rather than of what one *does* believe, still what one is justified in believing is a function of what one *does* grasp or understand.) Lest this be thought objectionable, it is worth pointing out that even adherents of View 1 are likely to take obligation to be subjective, in the present sense. (Whether this is in fact so will depend on how exactly they develop their view.) Moore himself, for example, takes obligation to be restricted to those acts that he calls 'voluntary' (in the sense that one would perform them if one chose).[68] Moreover, it seems clear that any plausible view of obligation must take it to be subjective, in the present sense. This is because whether one is a moral agent (in the sense of being capable of having moral obligations) is surely in part a function of one's mental capacities. Normal adults are moral agents, whereas infants and non-human animals are not, and this has at least in part to do with the difference in their mental capacities.

[68] Moore, *Ethics*, pp. 5 ff.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press

Views 1 and 2b would thus appear to be on a par with respect to whether moral obligation is objective or subjective, in the two senses just discussed. Perhaps there is a third sense that would serve to distinguish these views in this regard. For example, one might claim that, although View 1 should be understood to imply, as View 2b does, that an agent's state of mind is an 'enabling condition' of his being obligated to perform some act,[69] it does not imply, as View 2b does, that that act's being obligatory 'supervenes' at all on the agent's mental state. In the end, though, whether we say that overall moral obligation is 'objective' or 'subjective' is an issue of small moment. What matters is whether View 2b is true. If it is not, we must search for another alternative to View 1, for that view is certainly false.[70]

m_zimme2@uncg.edu

[69] Compare Jonathan Dancy, *Ethics without Principles* (Oxford, 2004), pp. 38 ff.

[70] Earlier versions of this article were presented to audiences at the following universities: Cambridge, Gothenburg, Lund, Oxford, Stockholm, and Uppsala. For comments on those or other occasions I am grateful to Robert Adams, Gustaf Arrhenius, Simon Blackburn, Johan Brännmark, John Broome, Åsa Carlson, Roger Crisp, Sven Danielsson, Daniel Elstein, Fred Feldman, Jonas Gren, Ross Harrison, Brad Hooker, Magnus Jiborn, Jens Johansson, Niklas Juth, Andreas Lind, Hans Mathlein, Ragnar Ohlsson, Jan Österberg, Tomasz Pol, Robert Pulvertaft, Wlodek Rabinowicz, Toni Rønnow-Rasmussen, Joakim Sandberg, Julian Savulescu, Torbjörn Tännsjö, and Ralph Wedgwood. I am especially indebted to Krister Bykvist, Erik Carlson, and Ingmar Persson. Finally, I thank the National Endowment for the Humanities for support during the writing of this article.

https://doi.org/10.1017/S0953820806002159 Published online by Cambridge University Press