

Luminosity Failure, Normative Guidance and the Principle ‘Ought-Implies-Can’

NICK HUGHES

University College Dublin

It is widely thought that moral obligations are necessarily guidance giving. This supposed fact has been put to service in defence of the ‘ought-implies-can’ principle according to which one cannot be morally obligated to do the impossible, since impossible-to-satisfy obligations would not give guidance. It is argued here that the supposed fact is no such thing; moral obligations are not necessarily guiding giving, and so the ‘guidance argument’ for ought-implies-can fails.

I. OUGHT-IMPLIES-CAN

Can you be morally obligated to do the impossible? Many philosophers think not. They subscribe to the following principle:

OIC: Necessarily, if a person is morally obligated to φ , then they can φ .

OIC is easiest understood when it is read as stating that if a person cannot perform some action φ , then they cannot be morally obligated to φ .¹ The demands of morality, in other words, never require one to do the impossible. Given that Asha cannot lift the ten-ton boulder in front of her, there is no question of her being morally obligated to, no matter what consequences might result from her inevitable failure. Suppose that Asha’s friend is trapped under the boulder with serious injuries and will die if she doesn’t lift it. Even so, *OIC* says, she isn’t obligated to lift it, because she can’t.²

OIC has played, and continues to play, a significant role in moral theory. To give just a few examples: it is widely taken to be a central pillar of Kant’s moral philosophy;³ it has been used to argue against

¹ This is the contrapositive of *OIC*.

² *OIC* is one of many ‘ought-implies-can’ principles, each of which is a variation on a logical theme, taking some sense of ‘ought’ and some sense of ‘can’ and saying for those senses that one ought to φ only if one can φ . The ‘ought’ in *OIC* is that of moral obligation. I have deliberately left the sense of ‘can’ unspecified. It won’t matter for my purposes, since everything I want to say in this article will apply to *OIC* for any interpretation of ‘can’.

³ Though there is some controversy about whether Kant should be interpreted as subscribing to it. See Robert Stern, ‘Does “Ought” Imply “Can”? And Did Kant Think It Does?’, *Utilitas* 16 (2004), pp. 42–61.

the possibility of moral dilemmas;⁴ it enters into debates about how to best formulate consequentialist theories;⁵ it has been appealed to in arguments against both objectivist and subjectivist conceptions of our moral obligations;⁶ and it is often invoked in debates about the compatibility, or otherwise, of determinism and free will,⁷ and debates about the relationship between moral obligations and alternative possibilities.⁸ This is only a snapshot. A complete catalogue of the debates and discussions in which the principle has been appealed to would be much longer.

Given its exalted status, one would expect *OIC* to be very well motivated, for if it is false then whole swathes of moral theory rest on a mistake. In fact, however, it is contentious at best. Whilst several arguments have been put forward purporting to motivate it, each has been subjected to trenchant criticism.⁹ In this article I will add another voice to the sceptical chorus. In particular, I will argue that one of the most widely endorsed arguments for the principle – what I will call the ‘guidance argument’ – should be considered a failure.¹⁰ The reason why may be surprising. I will argue that it is undermined by recent insights in epistemology about the limits of our knowledge.

⁴ See, in particular, the essays collected in Christopher Gowans, *Moral Dilemmas* (New York, 1987) and H. E. Mason, *Moral Dilemmas and Moral Theory* (Oxford, 1996).

⁵ See Elinor Mason, ‘Consequentialism and the “Ought Implies Can” Principle’, *American Philosophical Quarterly* 40 (2003), pp. 319–331, and Vuko Andric, ‘Objective Consequentialism and the Rationales of “Ought” Implies “Can”’, *Ratio* 29 (2015), pp. 1–16.

⁶ See Peter Graham, ‘“Ought” and Ability’, *Philosophical Review* 120 (2011), pp. 337–82; Mason, ‘Consequentialism’; Andric, ‘Objective Consequentialism’.

⁷ See Ishitiyaque Haji, *Moral Appraisability: Puzzles, Proposals, and Perplexities* (New York, 1998); Ishitiyaque Haji, ‘Moral Anchors and Control’, *Canadian Journal of Philosophy* 29 (1999), pp. 175–203; Ishitiyaque Haji, *Deontic Morality and Control* (New York, 2002).

⁸ See Peter van Inwagen, *An Essay on Free Will* (Oxford, 1983); David Widerker, ‘Frankfurt on “Ought Implies Can” and Alternative Possibilities’, *Analysis* 51 (1991), pp. 222–4; David Copp, ‘Defending the Principle of Alternate Possibilities: Blameworthiness and Moral Responsibility’, *Nous* 31 (1997), pp. 441–56; David Copp, ‘“Ought” Implies “Can”, Blameworthiness, and the Principle of Alternate Possibilities’, *Moral Responsibility and Alternative Possibilities: Essays on the Importance of Alternative Possibilities*, ed. Michael McKenna and David Widerker (Burlington, VT, 2003), pp. 265–99.

⁹ For a good overview of the broader debate, see Peter Vranas, ‘I Ought, Therefore I Can’, *Philosophical Studies* 136.3 (2007), pp. 351–84.

¹⁰ My view is that *OIC* is false, and so every argument employing it as a premise is unsound. But I won’t argue for that claim here. My ambition is only to show that the guidance argument in particular is a failure. This leaves open the possibility that the principle may be established on other grounds (though as I said, I’m sceptical).

II. THE GUIDANCE ARGUMENT

The guidance argument for *OIC* is advanced by Hare, Driver, Smith, Williams, Griffin, Copp, and Andric, amongst others.¹¹ It goes like this:

1. If you are morally obligated to φ , then you can be guided in your deliberation about whether or not to φ by this fact.
2. You can be guided in your deliberation about whether or not to φ by your moral obligation to φ only if you can φ .
3. Therefore, you are morally obligated to φ only if you can φ .¹²

Clearly the argument is valid: (3) follows straightforwardly from (1) and (2) by hypothetical syllogism. But what of the premises? Why should we accept them?

The thought behind premise (1) is that one of the central purposes of moral obligations, perhaps *the* central purpose, is to provide people with guidance (advice, instruction, etc.) about how to act. The foundational question of moral theory is, after all, and as it is often said, 'What should we *do*?'. A theory according to which you can be obligated to φ , yet cannot be guided in your deliberation about whether or not to φ by the obligation, must then be rejected on the grounds that it fails to live up to its *raison d'être*: to guide your deliberation about what to do when it comes to the choice between φ -ing and not φ -ing. In the course of advancing the guidance argument, Copp puts the idea pithily: "The point of moral requirements is to affect our decisions, and to lead us to do what is right, by being taken into account in our deliberation."¹³

The thought behind premise (2) is that when a moral theory requires someone to do what they are unable to do it is entirely useless at fulfilling this purpose. In deliberating about what to do, one only takes into account the available options. If you cannot φ , then being told

¹¹ See amongst others: R. M. Hare, *Reason and Freedom* (Oxford, 1963); Julia Driver, 'Promises, Obligations, and Abilities', *Philosophical Studies* 44 (1983), pp. 221–3; Holly Smith, 'Moral Realism, Moral Conflict, and Compound Acts', *Journal of Philosophy* 83.6 (1986), pp. 341–5; Bernard Williams, 'Internal Reasons and the Obscurity of Blame', *Making Sense of Humanity* (Cambridge, 1998); James Griffin, 'The Human Good and the Ambitions of Consequentialism', *The Good Life and the Human Good*, ed. E. F. Paul, F. D. Miller and Jeffrey Paul (Cambridge, 1992), pp. 118–32; Copp "Ought" Implies "Can"; Andric, 'Objective Consequentialism'.

¹² Usually the argument is not laid out precisely, step by step, but instead put forward in a more casual and offhand way. This is unfortunate, as it can make it difficult to see exactly how it is supposed to go. In 'Guidance, Obligations, and Ability: A Close Look at the Action Guidance Argument for Ought-Implies-Can', *Utilitas* (forthcoming), I argue that if it is to have any plausibility at all the argument must be understood as taking the form of (1)–(3).

¹³ Copp, "Ought" Implies "Can", p. 273.

that you ought to φ won't make the slightest bit of difference to your deliberation about whether or not to φ , for the simple reason that – if you are rational, at least – you won't deliberate about whether or not to do so in the first place.¹⁴ Sartre once memorably described an occasion on which he was approached by a student torn between leaving home to fight against fascism in Vichy France – a cause he rightly regarded as just – and staying home to care for his elderly mother, for whom he was the only source of solace and comfort in life. Which course, the student wanted to know, should he take? Had Sartre's response been 'Both. You should stay at home to care for your mother and leave home to fight the forces of fascism', the student would have no doubt, and quite reasonably, complained that this answer provided him with no guidance at all. He could not both leave home and stay home, and he wanted to know what he should do of the things he was actually able to do. The motivation for premise (2) is the thought that the same can be said about you, me, the student, Asha, and everyone else who seeks guidance.¹⁵

III. THE CONDITIONS OF GUIDANCE

At first glance the guidance argument appears to provide a compelling motivation for *OIC*. Each of its premises looks plausible, and the conclusion that *OIC* is true follows straightforwardly from them. Why do I think it fails, then? Because premise (1) is false. You can be morally obligated to φ even though you cannot be guided in your deliberation about whether or not to φ by this fact, or so I'll argue.¹⁶

To see why, first we need to look at the conditions which must be met in order for you to be guided by an obligation. We've already seen one such condition. You can only be guided in your deliberation about whether or not to φ by an obligation to φ if you are able to φ . It is this fact that the guidance argument appeals to. But that's not the only

¹⁴ Provided you know that you cannot φ , at least.

¹⁵ Jean-Paul Sartre, *Existentialism and Humanism* (London, 1948). Sartre's response wasn't 'both', by the way. It was that no moral theory can help the student; he must simply choose. Perhaps he wouldn't have accepted the guidance argument, then.

¹⁶ A referee has suggested that we should also reject premise (2) of the argument. If you are unable to φ , but don't know it, couldn't an obligation to φ guide you in trying to φ ? Perhaps. But those who endorse the guidance argument may maintain that deliberating about whether or not to try to φ is not the same thing as deliberating about whether or not to φ . I might know that I can't pick up a large boulder, but nevertheless decide to try to, because by trying and failing I will thereby demonstrate to my friend that it can't be done (see Jennifer Hornsby, 'Reasons for Trying', *Journal of Philosophical Research* 20 (1995), pp. 525–39). I've certainly deliberated about whether or not to try to pick up the boulder in this case, but have I deliberated about whether or not to pick up the boulder *simpliciter*? One reason to think not is that it would be irrational for me to deliberate about whether or not to pick it up, since I know that I can't, but it isn't irrational for me to deliberate about whether or not to try.

condition which must be met if you are to be guided. For suppose that you are obligated to φ , but not in a position to know about it.¹⁷ If so, then even if you are able to φ , and even if you do φ , it is hard to see how your decision to φ could have been *guided* by the obligation. Smith gives an illustrative example. If you order fish rather than pork at a restaurant, then you have acted in accordance with the Levitical law prohibiting consumption of any animal that 'parts the hoof but fails to chew the cud'. But if you are wholly ignorant of Levitical law, then you haven't been guided in your choice by it.¹⁸ The fact that you've done what Levitical law requires of you is nothing more than a coincidence.

Here's another example, which I will draw on later. Suppose that Asha is a doctor, and that she's obligated to give her patient an injection of morphine, because he's in extreme pain, and that's what she's obligated to do whenever he's in extreme pain. Suppose, however, that Asha doesn't know that she's subject to the obligation, because she doesn't know that the patient is in extreme pain (he's in such a bad way that there are no external signs of his inner torment). Even if there is morphine available, and Asha does in fact inject the patient with it, again, whilst she has done what she was obligated to do, she hasn't been *guided* in her behaviour by the obligation.¹⁹ As before, the fact that she has done what's required of her is nothing more than a coincidence.

What these cases show is that in order to be guided by an obligation to φ , you must not only be able to φ , you must also know that you are obligated to φ .²⁰ Both conditions must be met if you are to be guided. If one is satisfied, but the other isn't, then it is not as though you

¹⁷ Some readers might think that one cannot be obligated to φ but not in a position to know about it. If you are one of them, I ask you to hold that thought. I will address it in the next section.

¹⁸ Holly Smith, 'Using Moral Principles to Guide Decisions', *Philosophical Issues* 22 (2012), pp. 369–86. I don't mean to suggest that conforming to Levitical law is actually morally obligatory (nor, I assume, does Smith). It merely provides a useful example for the point I want to make.

¹⁹ We are interested here (and later) in the obligation to give the patient an injection of morphine on this occasion, rather than the more general conditional obligation to give the patient an injection of morphine if he's in extreme pain. The conditional obligation should be read as having narrow scope (that is, as having the logical form $A \rightarrow OB$, rather than $O(A \rightarrow B)$, where 'A' = the patient is in extreme pain, 'B' = the patient is given an injection of morphine and 'O' = it is obligated that), otherwise we cannot detach the obligation to give the patient an injection of morphine on this occasion from the conditional obligation. Thanks to a referee for pointing this out.

²⁰ I'm not saying anything new here. The observation that you can only be guided by an obligation to φ if you know about it has been made many times before. See amongst others: James Hudson, 'Subjectivization in Ethics', *American Philosophical Quarterly* 26.3 (1989), pp. 221–9; Allen Gibbard, *Wise Choices, Apt Feelings* (Oxford, 1990); Frank Jackson, 'Decision-Theoretic Consequentialism and the Nearest and Dearest Objection', *Ethics* 101(1991), pp. 461–82.

get some sort of partial guidance from the obligation. You get none whatsoever.²¹

IV. GUIDANCE FAILURE AND THE LIMITS OF KNOWLEDGE

So, if you are to be guided in your deliberation about whether or not to φ by an obligation to φ , you must both be able to φ and know that you are obligated to φ . Both conditions must be met, and if one or the other isn't, you get no guidance whatsoever. This is where the problem lies for the guidance argument. For there are good reasons to think that you can be obligated to φ even though you don't know it, and even though you are not even in a position to know it. If so, then there are possible circumstances in which you are obligated to φ but incapable of being guided in your deliberation about whether or not to φ by this fact. And if that's right, then premise (1) of the guidance argument is false and the argument must be rejected.

The reasons stem from a by-now-famous argument put forward by Williamson in *Knowledge and its Limits*. Williamson argues that no non-trivial condition is 'luminous', where luminosity is defined as follows:

Luminosity: A condition C is luminous if and only if whenever C obtains, one is in a position to know that C obtains.²²

Williamson makes the case for no non-trivial condition being luminous by first arguing that a condition which looks like an especially plausible candidate for being luminous – the condition of *feeling cold* – fails to satisfy *Luminosity*, and then generalising on the grounds that we can run the same argument for any non-trivial condition. Since Williamson's argument will be central to my case against the guidance argument, I will go through it in some detail.²³

²¹ Doubtless there are other conditions which must also be satisfied, but they won't matter for my argument.

²² The class of trivial conditions is comprised mainly of those that obtain in all or no possible worlds. Clearly there are non-trivial moral obligations in this sense. Asha may be obligated to inject her patient with morphine in this world, because the patient is in extreme pain, but there is a possible world in which Asha has no such obligation, because the patient isn't in pain at all. There may be trivial conditions that don't obtain in all or no possible worlds – the condition that one exists (under a first-personal mode of presentation), for instance – but as Williamson points out, if such conditions do exist they are curiosities at best. They won't affect my argument.

²³ Amia Srinivasan, 'Normativity without Cartesian Privilege', *Philosophical Issues* 25.1 (2015), pp. 273–99, also draws consequences for ethical theory from Williamson's anti-luminosity argument. Her goal is to show that if no non-trivial condition is luminous, then we must accept a tragic view of the normative realm, according to which it is ineliminably suffused with luck.

Roughly stated, the argument rests on two claims: that for a person's belief to be an item of knowledge it must be *safe from error*, and that for every non-trivial condition C there are *liminal cases pairs*: pairs of cases in which C obtains in one, but not the other, but which are so similar that the person cannot tell them apart. Since knowledge requires safety from error, the person only knows that C obtains if they could not have easily falsely believed that it obtained. But when it comes to liminal cases the person's true belief that C obtains is not safe: they could have easily falsely believed that it obtained, because there is a very similar case in which it does not obtain but – since the person cannot tell the two cases apart – they will be disposed to believe that it does. Thus, there is a case in which C obtains, but the person is not in a position to know that it obtains.

This can and should be put more precisely. Consider the following case:

Cold Morning: Asha wakes up at dawn feeling freezing, very slowly warms up, and feels hot by noon. Throughout the morning she is concentrating sufficiently hard on the question of whether she feels cold, such that if she is in a position to know that she feels cold, then she does indeed know. Asha's powers of discrimination are limited, and the change from her feeling cold to hot is so gradual that she is not aware of any change of feeling over one millisecond. Asha's confidence that she feels cold gradually diminishes, such that by noon she firmly believes that she no longer feels cold.²⁴

The more precise version of the argument goes like this.²⁵ Let t^0, \dots, t^n be a series of cases one millisecond apart from one another, let C be the condition that Asha feels cold, and let K(C) be the condition that Asha knows that she feels cold. Now assume that C is luminous. This gives us:

(LUM) If C obtains at t , then K(C) obtains at t

As already stated, knowledge requires safety. That is to say, one knows that P only if one could not have easily falsely believed that P. Williamson argues that the safety requirement motivates the following margin-for-error principle:

(MAR) If K(C) obtains at t , then C obtains at t^{+1}

²⁴ Adapted from Timothy Williamson, *Knowledge and its Limits* (Oxford, 2000).

²⁵ My presentation of the argument here draws on Selim Berker, 'Luminosity Regained', *Philosophers' Imprint* 8.2 (2008), pp. 1–22, and Amia Srinivasan, 'Are We Luminous?', *Philosophy and Phenomenological Research* 90.2 (2013), pp. 294–319.

The thought here is that if C doesn't obtain at t^{+1} , then one does not know that it obtains at t . This is because t and t^{+1} are so similar that one cannot tell them apart, and so if one believes that C obtains at t , one could have easily falsely believed that C obtains at t^{+1} . In that case one's belief that C obtains at t is not safe, and thus not an item of knowledge.

The *Cold Morning* case stipulates the following premises:

- (BEG) C obtains at t^0
- (END) C does not obtain at t^n

But (LUM), (MAR), (BEG) and (END) jointly lead to a contradiction, as follows:

1. By (BEG) C obtains at t^0
2. By (LUM) and (1), K(C) obtains at t^0
3. By (MAR) and (2), C obtains at t^{0+1}
4. Steps (1)–(3) are indefinitely repeatable²⁶
5. Therefore, C obtains at t^n
6. By (END) C does not obtain at t^n
7. (6) contradicts (5)

So, to avoid contradiction we must give up one of (LUM), (MAR), (BEG), or (END). (BEG) and (END) are true by stipulation, so that leaves us with (LUM) and (MAR). Since (MAR) is well motivated, we should give up (LUM). Feeling cold is not a luminous condition. Whilst Asha may usually know that she feels cold when she feels cold, she does not always know, because she doesn't know in liminal cases. And since one can construct cases with the same structure as *Cold Morning* for every non-trivial condition, we should conclude that no non-trivial condition is luminous.²⁷

For our purposes, the important condition is *being morally obligated to φ* . Let's suppose that if her patient *appears to her* to be in extreme pain, then Asha is obligated to give him an injection of morphine.²⁸ Here is a case analogous to *Cold Morning*:

²⁶ To see why, notice that if C obtains in t^{0+1} then by (LUM) K(C) obtains at t^{0+1} . And if K(C) obtains at t^{0+1} then by (MAR) C obtains at t^{0+2} . And if C obtains in t^{0+2} , then by (LUM) K(C) obtains at t^{0+2} . And if K(C) obtains at t^{0+2} then by (MAR) C obtains at t^{0+3} . . . and so on all the way to t^n .

²⁷ If you don't find this argument convincing, please hold that thought. I will address it in the next section.

²⁸ Again, the conditional obligation should be read as having narrow scope here, and our interest is in the detached obligation to give the patient an injection of morphine on a particular occasion (see n. 17).

Pain: At t^0 the patient wakes up from an operation. He appears to Asha to be in no pain whatsoever. Over the course of the day he appears to feel more and more pain, and by t^n , 12 hours later, he appears to be in extreme pain. Throughout the day Asha is concentrating sufficiently hard on the question of whether the patient appears to be in extreme pain, such that if she is in a position to know that he does, then she does indeed know that he does. Asha's powers of discrimination are limited, however, and the change from the patient appearing painless to appearing to be in extreme pain is so gradual that she is not aware of any change in how he appears to her, pain-wise, over one millisecond. Asha's confidence that he appears to be in pain gradually increases over time, such that at t^0 she firmly believes that he does not appear to be in extreme pain, and at t^n she firmly believes that he appears to be in extreme pain.

Earlier I described an obligation to give the patient an injection generated by him *in fact* being in extreme pain. Here I've described an obligation to give the patient an injection of morphine generated by him *appearing* to be in extreme pain. There is a reason for this change of focus. It is obvious that if Asha is obligated to give the patient an injection of morphine when he is in fact in extreme pain, then there can be cases in which one is obligated to φ but not in a position to know about it (again, the patient may be in such a bad way that there are no external signs of their inner torment). But some ethicists will be resistant to the idea that one can be morally obligated to φ but not in a position to know about it.²⁹ They will prefer a more internalist conception of our moral obligations. The would-be obligation for Asha to give the patient an injection of morphine generated by the appearance-to-her of him being in extreme pain is more internalist-friendly, and it is far less obvious that one can be subject to such an obligation yet not in a position to know about it. However, if the Williamsonian anti-luminosity argument is sound, then there will indeed be such cases, since we can reason in exactly the same way as in the *Cold Morning* case. Let t^0, \dots, t^n be a series of cases one millisecond apart from one another, and let us make the following two stipulations:

(BEG) The patient does not appear to Asha to be in extreme pain at t^0

(END) The patient appears to Asha to be in extreme pain at t^n

The argument goes as follows:

²⁹ For instance: Hudson, 'Subjectivization in Ethics'; Gibbard, *Wise Choices, Apt Feelings*; Jackson, 'Decision-Theoretic Consequentialism'.

1. By (BEG) the patient does not appear to Asha to be in extreme pain at t^0
2. By (LUM) and (1), Asha knows that the patient does not appear to her to be in extreme pain at t^0
3. By (MAR) and (2), the patient does not appear to Asha to be in extreme pain at t^{0+1}
4. Steps (1)–(3) are indefinitely repeatable
5. Therefore, the patient does not appear to Asha to be in extreme pain at t^n
6. By (END) the patient appears to Asha to be in extreme pain at t^n
7. (6) contradicts (5)

So, to avoid contradiction we must give up one of (LUM), (MAR), (BEG) or (END). As before, (BEG) and (END) are true by stipulation, so we must choose between (LUM) and (MAR). But again, (MAR) is well motivated. So we should give up (LUM). The condition that the patient appears to Asha to be in extreme pain is not luminous: even if she is usually in a position to know when it obtains, there are liminal cases in which it obtains but she is not in a position to know it. These are the cases in which the patient is in extreme pain that are so similar to cases in which the patient isn't in extreme pain that she cannot tell the two apart.

The conclusion we should draw from this is that one can be morally obligated to φ but not in a position to know about it. No matter how internalist we go with our conception of moral obligations, there will always be possible cases – the liminal cases – in which one is obligated to φ but not in a position to know it. You might not agree with the claim that a moral obligation for Asha to give a patient morphine can be generated by him appearing to her to be in extreme pain. One reason might be because you don't think that giving injections of morphine is the kind of thing that doctors can be morally obligated to do. If so, choose a different case – one that describes the kind of thing you think *can* be morally obligatory. It won't make a difference, as we will be able to run the argument with that case as well. Another reason might instead be because you don't think that appearances matter when it comes to moral obligations. Perhaps it is beliefs about how things are, or what the available evidence indicates about how things are, or something like that. But again, this won't make a difference – the argument will still go, because *no* non-trivial condition is luminous. The conclusion is unavoidable.

Where does this leave us? As I see it, with a conclusive reason to reject premise (1) of the guidance argument. As we saw earlier, if you are to be guided in your deliberation by an obligation to φ , you

must both be able to φ and know that you are obligated to φ . Given luminosity failure, there are cases in which you are obligated to φ but don't know it. It follows that there are cases in which you are obligated to φ but cannot be guided in your deliberation about whether or not to φ by the obligation. But this directly contradicts premise (1) of the guidance argument, which states that if you are morally obligated to φ , then you can be guided in your deliberation about whether or not to φ by this fact. So premise (1) is false, and the argument must be rejected. Whether or not *OIC* is in fact true, then, it isn't motivated by the guidance argument.

This might seem complicated, but the basic idea is simple. It boils down to this: the whole point of the guidance argument is that if *OIC* is false, then there will be cases in which you cannot be guided by a moral obligation. But, the guidance argument goes, there are no such cases, as morality is necessarily guidance giving. So *OIC* must be true. However, if no non-trivial condition is luminous, then there are such cases. And it follows that there is no motivation for *OIC* from the guidance argument.³⁰

V. OBJECTIONS, REVISIONS AND REPLIES

Or so I claim, at least. I am, however, aware that there are potential objections to this argument. Here I will address what I take to be the three most conspicuous of them.

V.1. Objection: the Williamsonian argument is unsound

First, my argument obviously stands or falls with the Williamsonian anti-luminosity thesis. You might not find Williamson's argument convincing, in which case you will presumably want to reject my argument too. I don't propose to defend the anti-luminosity argument here, since to do so would take us too far afield, and it has been ably defended by others. What I will say, however, is that it is quite widely accepted amongst epistemologists. That's not to say that it hasn't been subjected to criticism. What philosophical argument hasn't? The critics fall mainly into one of two camps. The first camp claims that knowledge

³⁰ It occurs to me that this argument might seem to rely on a sleight of hand. I've argued that there is one way in which you can fail to be guided by a would-be obligation – by not knowing about it – and I have used this point to argue that for a different way in which you can fail to be guided – by not being able to fulfil the would-be obligation – the fact of this failure does not undermine the possibility that the obligation is nevertheless a genuine one. Is this a legitimate form of argument? Indeed it is. If you are in a situation in which you *already* can't be guided by an obligation, because you don't know about it, then there is simply no motivation for insisting that you must be able to fulfil the obligation because otherwise you can't be guided by it. It's too late. You can't be guided anyway. That ship has already sailed.

doesn't require safety, and so that the margin-for-error principle (MAR) is not well motivated,³¹ or accepts the safety condition, but argue that it doesn't motivate (MAR).³² The second camp claims that the argument implicitly relies on a tolerance principle and is thereby soritical.³³ My view is that close inspection of the argument reveals that both camps are mistaken, and that Srinivasan has shown as much.³⁴ Nevertheless, if you are sceptical you are welcome to read my argument as disjunctive: either the guidance argument fails, or moral obligations are luminous. Given that the guidance argument and the anti-luminosity argument are both widely accepted, the news that they are incompatible should be interesting in its own right.³⁵

V.2. *Objection: you don't need to know to be guided*

A second objection accepts the anti-luminosity argument, but denies that you must know that you are subject to an obligation in order to be guided by it. One might think that an obligation can guide your deliberation even if you don't know that you are subject to it, provided that you believe that you are, or provided that you rationally believe that you are, or provided that it is probable to some degree n on your evidence that you are. If so, then one may agree with the claim that you can be subject to an obligation but not in a position to know about it, but nevertheless maintain that premise (1) of the guidance argument stands firm, because knowledge isn't necessary for guidance.

This response has two flaws, both of which seem to me to be fatal. The first is this. Notice to begin with that in order for your decision to

³¹ See amongst others: Ram Neta and Guy Rohrbaugh, 'Luminosity and the Safety of Knowledge', *Pacific Philosophical Quarterly* 85 (2004), pp. 296–406, and Earl Conee, 'The Comforts of Home', *Philosophy and Phenomenological Research* 70 (2005), pp. 444–51.

³² Brian Weatherston, 'Luminous Margins', *Australasian Journal of Philosophy* 82 (2004), pp. 373–83; Jonathan Vogel, 'Luminosity and Indiscriminability', *Philosophical Perspectives* 24.1 (2010), pp. 547–72.

³³ Wai-hung Wong, 'What Williamson's Anti-Luminosity Argument Really Is', *Pacific Philosophical Quarterly* 89 (2008), pp. 536–43; Berker, 'Luminosity Regained'. A tolerance principle is a principle of the form 'If C obtains in case n , then C obtains in case $n^{\pm 1}$ '. It is by relying on the apparent plausibility of some tolerance principles that soritical arguments create a paradox.

³⁴ Srinivasan, 'Are We Luminous?'.

³⁵ How pervasive is the problem? If, as a referee has suggested, we say that all obligations *except* the ones that depend on vague predicates are action guiding, will that help? It won't, for two reasons. First, as Williamson himself points out, the anti-luminosity argument doesn't only apply to conditions articulated with vague predicates, it applies equally to conditions articulated with sharp predicates (Williamson, *Knowledge and its Limits*). Second, even if that wasn't true, the most that those who put forward the guidance argument would be able to conclude, on this construal of it, is that *OIC* holds in cases where the obligation doesn't depend on vague predicates. But if so, it isn't fit to play the roles in moral theory described in section I, which depend on its holding in every case.

φ to be guided by an obligation to φ , the fact that you are obligated to φ must be among your motivating reasons for deciding to φ (even if it is not your only motivating reason). A person who φ s, but does so for reasons unrelated to the fact that they were obligated to, hasn't been *guided* by the obligation any more than someone who has no idea they are obligated to φ but happens to do so anyway. I may order fish rather than pork, knowing that I thereby conform to the demands of Levitical law. But I may not give a damn what Levitical law demands of me. Its demands may not be amongst the reasons for which I chose the fish over the pork and I might have chosen the pork just as easily, had the fish seemed less appealing. In that case I haven't been guided in my choice by the (supposed) obligation not to consume any animal that parts the hoof but fails to chew the cud. I've conformed to it, but just as with conforming to an obligation one is ignorant of, this is nothing more than a coincidence.³⁶

So in order for your decision to φ to be guided by an obligation to φ , the fact that you are obligated to φ must be among your motivating reasons for deciding to φ .³⁷ The problem for the response we are currently considering is that it is plausible that one's motivating reason for φ -ing can be P only if one knows that P.³⁸ The following case, adapted from Hyman,³⁹ supports this view:

³⁶ This isn't just true of obligations. In general, in order for my φ -ing to have been guided by some fact P, P must be amongst my motivating reasons for φ -ing. For instance, for my decision to take an umbrella to have been guided by the fact that it's raining, the fact that it's raining must be amongst my motivating reasons for deciding to take an umbrella.

³⁷ It is important to pay attention to the fact that we're talking about *motivating* reasons. When explaining a person's φ -ing, we can cite reasons that are not amongst their motivating reasons. For instance, the fact that Bill has a brain tumour may be a reason why he just lashed out (the tumour has changed his personality and made him prone to violence). It is an explanatory reason, but it need not have been amongst his motivating reasons, nor need it have guided his decision to φ – Bill may be totally unaware of the tumour. A referee has suggested that in cases in which a person is guided to φ by an obligation to φ , we can give a full *explanation* of their action whilst appealing only to their psychological states (beliefs, desires, etc.). I disagree. In appealing only to their psychological states, the explanation will miss out the fact that they φ 'ed not (merely) because they believed that they were obligated to, but because they were *in fact* obligated to.

³⁸ This claim is defended by: Peter Unger, *Ignorance* (Oxford, 1975); John Hyman, 'How Knowledge Works', *Philosophical Quarterly* 49.197 (1999), pp. 433–51; John Hyman, 'Knowledge and Evidence', *Mind* 115.460 (2006), pp. 891–916; John Hyman, *Action, Knowledge, and Will* (Oxford, 2015); Williamson, *Knowledge and its Limits*; Jennifer Hornsby, 'Knowledge in Action', *Action in Context*, ed. A. Leist (Berlin, 2007), pp. 285–302; Clayton Littlejohn, 'The Russellian Retreat', *Proceedings of the Aristotelian Society* 113.3 (2013), pp. 293–320. I argued against it in Nick Hughes, 'Is Knowledge the Ability to φ for the Reason that P?', *Episteme* 4.11 (2014), pp. 457–62 (see also Dustin Locke, 'Knowledge, Explanation, and Motivating Reasons', *American Philosophical Quarterly* 52.3 (2015), pp. 215–32). I have since changed my mind.

³⁹ Hyman, 'How Knowledge Works'.

Tennis: Asha is watching the television on a June afternoon. It is Wimbledon Men's Finals Day, and the television shows Federer beating Nadal; the score is two sets to love and match point to Federer in the third. Federer wins the point, and Asha forms the belief that Federer has won the match. However, unbeknownst to Asha, the cameras at Wimbledon have ceased to function, and the television is showing a recording of last year's final. But while it does so Federer has, unbeknownst to Asha, just finished a repeat of last year's slaughter.

Asha's belief that Nadal has lost this year's championship is rational and true, but not an item of knowledge because it is 'Gettiered'. Let's suppose that Asha is a Nadal fan, and that she is moved to write him a commiseratory letter. Why has she written the letter (i.e. what is her motivating reason for writing it)? As Hyman points out, the answer cannot be because Nadal has lost this year's championship, since Asha isn't even aware of that fact. The fact that Nadal has lost this year's championship appears nowhere in a catalogue of her motivations for writing the letter. At most we might say that she is motivated to write it because she *believes* that Nadal has lost this year's championship.⁴⁰

Why is Asha's belief that Nadal has lost this year's championship, rather than the fact that he has, her reason for writing the letter, despite the fact that her belief is both rational and true? The obvious answer is: because it is not an item of knowledge, and one can φ for the reason that P only if one knows that P.

If that's right, as I think it is, and if you have only been guided by an obligation to φ if the fact that you are obligated to is amongst your reasons for φ -ing, then it follows that even if you rationally (and

⁴⁰ As a referee has pointed out, this stands in need of further clarification. Asha's belief isn't her motivating reason if by this we mean that she treats the fact that she believes that Nadal has lost to be a reason to write him a letter. After all, it is not as though she'll think to herself: 'I believe that Nadal has lost the match, and I want to write him a letter if he has lost, so I'll write him a letter'. Rather, she'll just think 'Nadal has lost the match, and I want to write him a letter if he's lost, so I'll write him a letter'. So, what *is* her motivating reason then? One popular answer is that it is the content of her belief, which she takes to be a fact, rather than the fact that she believes it (see, amongst others, Jonathan Dancy, *Practical Reality* (Oxford, 2000); Jennifer Hornsby, 'A Disjunctive Conception of Acting for Reasons', *Disjunctivism: Perception, Action, Knowledge*, ed. Adrian Haddock and Fiona MacPherson (Oxford, 2008), pp. 244–62; Mark Schroeder, 'Having Reasons', *Philosophical Studies* 139 (2008), pp. 57–71). Another, more radical, answer is that she had no motivating reason at all, only the illusion of one (Maria Alvarez, *Kinds of Reasons: An Essay on the Philosophy of Action* (Oxford, 2010). I don't find either of these responses fully satisfying, but going into the issue further here would take us too far afield. Hereafter I will continue to speak of her belief being her motivation, but only for the sake of convenience; nothing much should be read into it.

truly) believe that you are obligated to φ you cannot be guided in your deliberation about whether or not to φ by the obligation if you don't know about it.

To illustrate the point, let's suppose that before the match started Asha bet her sister £50 that Nadal would emerge as this year's champion. Since she's lost the bet, she's obligated to pay up, which she does.⁴¹ But the *fact* that she is obligated to pay up is not her motivating reason for doing so any more than the fact that Nadal has lost is the reason for which she has written him a letter. The *fact* that she is obligated to do so appears nowhere in a catalogue of her motivations for paying up, after all. Her reason for paying up is at best that she *believes* she is obligated to so. Likewise, her decision to pay up wasn't guided by the *fact* that she was obligated to do so. It was at most her *belief* that she was obligated to do so, rather than the obligation *itself* (of which she was ignorant) that guided her decision.

So, putting together the idea that (a) you have been guided in your decision to φ by an obligation to φ only if the fact that you were obligated to φ is among the reasons for which you did so, and the idea that (b) P can be your reason for doing something only if you know that P, yields the conclusion that (c) you cannot be guided by an obligation if you don't know that you are subject to it.

If that's right, then the claim that rational belief is sufficient for guidance is false. Moreover, the same point applies, *a fortiori*, to the claim that you can be guided by an obligation provided that you believe that you are subject to it, and the claim that you can be guided by an obligation provided that it is probable on your evidence to some degree n that you are subject to it. It isn't the obligation itself doing the work in these circumstances, it's your belief that you are subject (or probably subject) to it.⁴²

⁴¹ Once again, the conditional obligation for Asha to pay up if she's lost a bet should be read as having narrow scope, and our interest is in the detached obligation (see n. 17).

⁴² A referee has suggested that you might be *indirectly* guided by an obligation to φ even if you don't know that you're obligated to φ . You might worry that, for all you know, you are subject to the obligation, and think to yourself 'better safe than sorry'. Has the obligation thereby (indirectly) guided your decision to φ ? My answer won't come as a surprise: no, it is at most your belief or knowledge that *you might be subject to the obligation* that has guided you here, not the obligation *itself*. Compare: for all I know China will top the medal table at the next Olympics. This fact about my epistemic position may guide my behaviour by, say, leading me not to accept certain bets against it happening (better safe than sorry). But it surely cannot be the case that the fact that China will top the medal table has guided my decision. For one thing, on some theories of facts, there is at present simply no such fact. But even if there is, it hasn't guided me, *even indirectly*. If it had, then supposing that China does indeed top the medal table, in 2020 I will be able to claim legitimately that the fact that they did (or would) was among my reasons for not betting against them in 2017. But that would be a ludicrous claim.

Those who subscribe to psychologism about motivating reasons won't accept this argument.⁴³ They maintain that it is never worldly facts that motivate people to act, but only ever their mental states. So while they'll agree that the fact that Asha is obligated to pay up isn't amongst her motivating reasons for doing so, they'll balk at the suggestion that it could have been, if only she hadn't been Gettiered. They'll insist that it still would have been her belief that motivated her even had she known that she was subject to the obligation. Might those who endorse the guidance argument appeal to psychologism at this juncture?

One obvious worry about this strategy is the tension between psychologism and our everyday practice of providing reasons. If you ask me why I'm leaving the party early, I may say it's because I have a meeting first thing tomorrow, and the fact that I have a meeting first thing certainly isn't a mental state of mine.⁴⁴ But even putting aside the demerits of psychologism itself, those who endorse the guidance argument are in no position to appeal to it in the first place. According to the psychologically inclined, the fact that I'm obligated to φ is not, and can never be, my motivating reason for φ -ing, for the simple reason that obligations aren't mental states. As I have just argued, my decision to φ has been guided by x (be it a fact, mental state, or whatever) only if x is among my motivating reasons for φ -ing. When we combine these two points, we get the result that I can only be guided in my decision to φ by my mental states. Since obligations aren't mental states, it follows that I can never be guided to φ by an obligation to φ . This is where the problem lies for those who endorse the guidance argument. For recall that their motivation for premise (1) of the argument is the idea that the very *point* of obligations is to guide our deliberation. If they were to embrace psychologism, then, they would have to conclude that all moral obligations are pointless. This is hardly something they're likely to be happy with. Worse still, since premise (2) of their argument says that if one cannot be guided by a would-be obligation, then it isn't a genuine obligation at all, they will also be forced to say that there are no moral obligations in the first place. Suffice to say, that's not the conclusion they're after.⁴⁵

⁴³ For example: Donald Davidson, 'Actions, Reasons, and Causes', *Journal of Philosophy* 60.23 (1963), pp. 685–700, and Michael Smith, *The Moral Problem* (Oxford, 1994). Thanks to an anonymous referee for pointing this out, and for encouraging me to think about how psychologism and the guidance argument interact.

⁴⁴ Dancy, in *Practical Reality*, is one of many to press this objection against psychologism.

⁴⁵ Might those who go in for the guidance argument somehow revise it in order to avoid this problem whilst nevertheless maintaining that you don't need to know to be guided? I find it hard to see how they could. But at any rate, even if it can be done, it is incumbent on *them* to show how.

All of the above seems to me to be a good reason to dismiss the response (which, remember, was that you can be guided by an obligation even if you don't know that you are subject to it). But there is another problem, which I think ought to dissuade even those who are so far unconvinced. Williamson has argued – again convincingly in my view – that the argument for the conclusion that no non-trivial condition is luminous can be extended to show that you can know that P even though you are not in a position to rationally believe that you know that P , and indeed, even though it is arbitrarily improbable short of 0 on your evidence that you know that P .⁴⁶ And as Hawthorne and Srinivasan have pointed out,⁴⁷ there is no reason to think that this argument won't likewise extend to all non-trivial conditions. If that's right, it follows that you can be obligated to φ even though it is not rational for you to believe that you are obligated to φ , and even though it is not probable to degree n on your evidence that you are obligated to φ (at least, if n has a value greater than 0). But if so, there will still be cases in which you are obligated to φ but unable to be guided by the obligation, even on these weaker interpretations of what is required for guidance. So premise (1) remains false.⁴⁸

V.3. *Objection: the guidance argument can be revised*

Finally, even if premise (1) of the guidance argument is false, it might be thought that a suitable replacement for it can be found that is not vulnerable to my argument, but still captures the spirit of the idea that guidance considerations motivate *OIC*. Specifically, one might think that even if it is inevitable that there will be some cases where you are obligated to φ but cannot be guided in your deliberation by this fact, the more often a moral theory is capable of giving guidance, the better. If so, then it might be argued that we should accept *OIC* because moral theories accepting it will give guidance more often than those rejecting it.

There are two ways of interpreting this argument. On the first interpretation the claim is that if moral theory *A* is guidance giving more often than moral theory *B*, this gives us a conclusive reason to

⁴⁶ It is difficult to briefly summarize Williamson's argument for these claims, and I won't attempt to do so here. The reader should consult Timothy Williamson, 'Very Improbable Knowing', *Erkenntnis* 79.5 (2014), pp. 971–99.

⁴⁷ John Hawthorne and Amia Srinivasan, 'Disagreement without Transparency: Some Bleak Thoughts', *The Epistemology of Disagreement*, ed. David Christensen and Jennifer Lackey (Oxford, 2013), pp. 9–30.

⁴⁸ I should note that this second reply doesn't apply to the claim that one can be guided by an obligation if one believes that one is subject to it, since the Williamsonian argument cannot be extended to show that one can be subject to an obligation even though one is not in a position to believe that one is subject to it. I am content to rely on my first reply to that claim here.

adopt *A* over *B*. On the second interpretation, the claim is that if theory *A* is guidance giving more often than theory *B*, this gives us a reason to adopt *A* over *B*, all other things being equal. Let's take these in turn.

I am unconvinced by the argument on the first interpretation, because the principle driving it will lead us to adopt absurdly subjectivist theories of moral obligation. The claim is that theory *A* is always superior to theory *B* if *A* gives guidance more often than *B*. But that cannot be right. A moral theory according to which you are obligated to φ only if you feel like φ -ing will be capable of giving guidance very often. You only need to ask yourself: do I feel like φ -ing? And whilst you are not always in a position to know what you feel like doing (due to luminosity failure) you are plausibly very often in a position to know. This theory will give guidance more often than, say, a theory according to which you are obligated to φ if your evidence indicates that φ -ing would be best (for some specified interpretation of 'best'), since, if you are normal, you will make mistakes about what your evidence indicates more often than you will make mistakes about what you feel like doing. If, as this argument claims, the fact that theory *A* is more often guiding than theory *B* is a conclusive reason to prefer *A* over *B*, then it would follow that we have a conclusive reason to prefer the theory claiming that you are obligated to φ only if you feel like φ -ing over the theory claiming that you are obligated to φ if your evidence indicates that φ -ing would be best. But that's absurd. The latter theory is clearly superior to the former (even if it isn't the theory we should ultimately end up adopting). The former is far too lax in the demands it makes on agents. There is more to morality than doing whatever you fancy.

So we should reject the argument on the first interpretation. The principle on which it relies is a bad one. But what about the second interpretation? Here the claim is much weaker: provided theory *A* and theory *B* are equally good in all other respects, *A* should be preferred to *B* if *A* gives guidance more often than *B*. And *A* will give guidance more often than *B* if *A* accepts *OIC* and *B* rejects it.

I have some sympathy with this suggestion. But it falls far short of giving us anything like a decisive (or even especially strong) reason to accept *OIC*. For it is a wide-open question whether other things are indeed equal when it comes to theories that accept *OIC* versus theories that reject it.⁴⁹ *OIC* has some pull, but it also comes with costs. For instance, many ethicists have thought that a realistic moral theory

⁴⁹ It is also questionable whether moral theories which accept *OIC* actually do better when it comes to giving guidance than those which reject it. See Brian Talbot, 'The Best Argument for "Ought Implies Can" Is a Better Argument against "Ought Implies Can"', *Ergo* 3.14 (2016), pp. 377–402, for an argument to this effect.

must allow for the possibility of tragic dilemmas in which one faces conflicting, jointly unsatisfiable, moral obligations.⁵⁰ If that's right, then it puts pressure on us to give up *OIC*, since it is a feature of tragic dilemmas that one is condemned to violate at least one obligation, whatever one does.⁵¹ Second, *OIC* is in tension with some seemingly platitudinous moral statements, such that one should never torture people for fun.⁵² If, due to some kind of psychological compulsion, a person simply cannot prevent themselves from torturing people for fun, and *OIC* is true, it follows that they are not morally obligated to refrain from doing so. That's a hard claim to countenance.⁵³

I don't want to take a stand here on whether these examples give us good reasons to reject *OIC*. They have, of course, been the subject of considerable debate. My point, rather, is this: putting guidance considerations aside, it is at best a matter of ongoing controversy whether *OIC* will feature in our best moral theories. Introducing *OIC* may give a theory a 'boost' on the grounds that it will thereby become better at giving us useful guidance, but it remains to be seen whether that boost will tip the dial in favour of a theory embracing *OIC* over one that rejects it, for we first need to take a great many other factors into account. As a result, there is no decisive argument for *OIC* to be found emerging from the idea that, *ceteris paribus*, moral theory *A* is better than *B* if *A* is more guiding than *B*. We're just not in a position yet to know how things will turn out. But even if it turns out that guidance considerations do tip the dial in favour of *OIC*, they will only have played a small role in the overall argument for the principle. There is no direct route, then, to establishing *OIC* from the idea that the point of a moral theory is to provide us with useful guidance about how to act.⁵⁴

nickhowellhughes@gmail.com

⁵⁰ See, amongst others, Bernard Williams, 'Ethical Consistency', *Proceedings of the Aristotelian Society, Supplementary Volumes* 39 (1965), pp. 103–38; Bas van Fraassen, 'Values and the Heart's Command', *Journal of Philosophy* 70.1 (1973), pp. 5–19; Ruth Barcan Marcus, 'Moral Dilemmas and Consistency', *Journal of Philosophy* 77.3 (1980), pp. 121–36; Martha Nussbaum, *The Fragility of Goodness* (Cambridge, 1986).

⁵¹ I should note though that there is a way of reconciling *OIC* with the possibility of tragic dilemmas: by rejecting the agglomeration principle, according to which if one ought to φ and one ought to ψ , then one ought to perform the complex action φ & ψ (see Williams, 'Ethical Consistency', and Hughes, 'Dilemmic Epistemology', *Synthese* (forthcoming)). But taking this approach comes with costs of its own. *OIC* is at least in tension with the idea that there can be tragic dilemmas, then.

⁵² See, amongst others, John Kekes, 'Ought Implies Can and Two Kinds of Morality', *The Philosophical Quarterly* 34.137 (1984), pp. 459–67.

⁵³ Here I have given two examples of problems for *OIC*, but more could be given.

⁵⁴ Thanks to John Hyman, Jennifer Hornsby, an audience at University College Dublin and two anonymous referees for this journal, for helpful discussion and feedback.