

GUIDANCE, OBLIGATIONS, AND ABILITY

A Close Look At The Action Guidance Argument For Ought-Implies-Can

*Nick Hughes, UCD
Forthcoming in Utilitas*

§0. Introduction

The principle that one can be morally obligated to ϕ only if one can ϕ is widely accepted. It has also played, and continues to play, a major role in moral theory. For example, it has important implications for how we should best formulate consequentialism and utilitarianism;¹ it is commonly understood as being a central pillar of Kant's moral philosophy;² various authors have argued from it to the conclusion that moral dilemmas are impossible;³ it is frequently invoked in debates about the compatibility of determinism and free will and debates about the relationship between moral obligations and alternative possibilities;⁴ and it has been claimed to be in tension with both objectivist and subjectivist conceptions of our moral obligations.⁵

¹ See Mason (2003) and Andric (2015).

² At least, as Kant is normally interpreted. There is some controversy here. See Stern (2004) for more details.

³ See in particular the essays collected in Gowans (1987) and Mason (1996).

⁴ On determinism and free will see Haji (1998, 1999, 2002). On alternative possibilities see van Inwagen (1983), Widerker (1991) and Copp (1997, 2003).

⁵ For an appeal to the principle to argue against subjectivism, see Graham (2011). Against objectivism (objective utilitarianism in particular), see again Mason (2003) and Andric (2015).

In light of the principle's significance for moral theory, one would expect it to be well motivated. But is it? Certainly plenty of arguments have been put forward in its favour. My interest here is only in one particular, frequently invoked, argument for it. We can call it 'the action guidance argument'. What I aim to show is that *this* particular argument fails to establish the principle. I am motivated to do so, in part, by my conviction that the argument has, to date, not been stated with sufficient precision and clarity, and that once it is brought into sharp relief it can be shown to be defective.

In §1 I give a precise statement of how I think the argument must go in order for it to be plausible. In §2 I present what I take to be a fatal problem with the argument, so understood. In §3 I consider how those who endorse the argument might respond to the problem, and argue that none of the responses works. In §4 I offer an alternative proposal for how to think about action guidance and moral obligations. In §5 I conclude.

§1. What Is The Action Guidance Argument?

The action guidance argument is put forward by, among others, Hare (1963), Driver (1983), Williams (1989), Griffin (1992), Copp (2003), and Andric (2015). It goes like this:

1. Necessarily, if you are morally obligated to ϕ , then you can be guided in your deliberation about how to act by this fact
2. Necessarily, you can be guided in your deliberation about how to act by the fact of your moral obligation to ϕ only if you can ϕ
3. Therefore, necessarily, you are morally obligated to ϕ only if you can ϕ

What exactly do those who put forward the argument have in mind when they say that, necessarily, if you are morally obligated to ϕ , then you can be guided in your deliberation about how to act by this fact? The question is a pressing one, since, as many others have pointed out (Martin 2009, Jay 2012, Stern 2004, among others), there are ways in which you can be guided in deliberation about how to act by a would-be obligation to ϕ even when you cannot ϕ . To give just one example, learning that you have an obligation to keep a promise to a friend can

guide your deliberation about whether or not to apologize if you fail to keep the promise. But an obligation to keep your promise can provide *this* guidance even if you are unable to fulfill the obligation, for the simple reason that you can still use the fact of your obligation to guide your deliberation about whether or not to apologize even after you learn that you will, inevitably, fail to fulfill it. Just because you can't ϕ , that doesn't mean you can't apologize for not ϕ -ing.

Evidently action guiders, as I'll call those who put forward the argument, must not have this kind of deliberative guidance in mind when they offer up premise (1). If they did, the argument wouldn't be remotely plausible. What is the idea then? Unfortunately it is rather difficult to find a clear and precise answer to this question in the literature. All too often those who endorse it put the argument forward in a rather casual and offhand way, without much elaboration, which makes it hard to determine just what is being proposed. Obviously this is unsatisfactory. It is impossible to evaluate the argument if we don't know exactly what its premises state.

Nevertheless, the situation is not hopeless. We can make progress on figuring out what action guiders must have in mind when we observe that an adequate answer to our question must satisfy two desiderata:

- It must be plausible that the proposed kind of guidance is given by an obligation to ϕ simply in virtue of it being an obligation to ϕ
- It must be plausible that the proposed kind of guidance necessarily cannot be given by an obligation to ϕ when you cannot ϕ

The first desideratum must be satisfied if premise (1) of the argument is to be well motivated. (1) Says that *necessarily*, if you are morally obligated to ϕ , then you can be guided in your deliberation about how to act by this fact. The relevant sense of necessity is presumably metaphysical. Were it not plausible that the proposed kind of guidance is given by an obligation to ϕ *simply in virtue of it being an obligation to ϕ* , then we would be left wondering just why we should think that an obligation to ϕ should give this kind of guidance *as a matter of metaphysical necessity*. The second desideratum must be satisfied if premise (2) of the argument is to be well motivated. If the proposed kind of guidance could be given by an obligation to ϕ even when you are unable to ϕ , then there would obviously be no reason to think that your deliberation can be guided in the relevant way by the fact of your obligation to ϕ only if you can ϕ .

So, what kind of guidance must action guiders have in mind if it satisfies both of these desiderata? The answer, I think, must be: *guidance in one's deliberation about whether or not to ϕ* . This fits our two desiderata rather well, and nothing

else appears to. It is plausible that an obligation to ϕ gives you guidance in your deliberation about whether or not to ϕ simply in virtue of being an obligation to ϕ because the very information conveyed by an obligation to ϕ , in virtue of it being an obligation to ϕ , is that you ought (indeed, must) ϕ rather than not- ϕ . And it is quite natural to think of this as a kind of guidance for one's deliberation about whether or not to ϕ . So the first desideratum appears to be satisfied. It is also plausible that an obligation to ϕ cannot guide your deliberation about whether or not to ϕ if you cannot ϕ , because if you cannot ϕ (or at least, if you *know* that you cannot ϕ) then, if you are rational, you will not deliberate about whether or not to ϕ in the first place, as such deliberation would be pointless. So an obligation to ϕ could not play the role of guiding your deliberation about whether or not to ϕ when you cannot ϕ for the simple reason that there would be no deliberation for it to guide in the first place. So the second desideratum appears to be satisfied as well.

This kind of deliberative guidance seems to satisfy the two desiderata well, then. But is there another kind of deliberative guidance that might also satisfy them? Whilst I have no proof that there isn't, I think it is highly doubtful, for two reasons. Firstly, the only information conveyed by an obligation to ϕ *simply in virtue of it being an obligation to ϕ* is that you must ϕ rather than not- ϕ . For this reason it is very hard to see what other kind of guidance an obligation to ϕ might be expected to give *as a matter of metaphysical necessity*. Secondly, when it comes to deliberation the only thing that is *necessarily* ruled out by your not being able to ϕ is (rational) deliberation about whether or not to ϕ . So it is likewise hard to see what other kind of deliberative guidance might motivate premise (2).

In that case, when the action guidance argument is fully spelled out it goes like this:

4. Necessarily, if you are morally obligated to ϕ , then you can be guided in your deliberation about whether or not to ϕ by this fact
5. Necessarily, you can be guided in your deliberation about whether or not to ϕ by the fact of your moral obligation to ϕ only if you can ϕ
6. Therefore, necessarily, you are morally obligated to ϕ only if you can ϕ

As I said, it seems to me that this is the most plausible way of understanding the action guidance argument. Is it convincing, so understood? It is not. As I'll show in the next section, it has unacceptable consequences.

§II. The Problem

The problem is this. In so far as it is true that if you are rational you will not deliberate about whether or not to ϕ when you know that you cannot ϕ , it is also true that (again, if you are rational) you will not deliberate about whether or not to ϕ when you know that you cannot *but* ϕ (i.e. when you cannot not- ϕ .) So just as there will be no deliberation for an obligation to ϕ to guide in cases where you cannot ϕ , there will likewise be no deliberation for an obligation to ϕ to guide in cases where you cannot but ϕ . The action guidance argument urges us to take the fact that a would-be obligation to ϕ could not play the relevant guiding role in cases where you cannot ϕ as a reason to think that you cannot be obligated to do what you are unable to do. But if we go along with that then we must, in order to avoid the charge of brazen inconsistency, also conclude that you cannot be obligated to do what you are unable *not* to do. And moreover, given the duality principle of deontic logic, according to which one is not obligated to ϕ if and only if one is permitted to not- ϕ ($\neg O \phi \leftrightarrow P \neg \phi$), we must also conclude that you are *permitted* to not- ϕ if you are unable to not- ϕ .⁶ That, however, is deeply implausible. To see why, consider the following case:

Football: I am at the home game of a football team who recently lost a player in a plane crash. The crowd is asked to observe two minutes silence in honor of his memory. Fortunately it is easy for me to comply with this request, and in fact it is impossible for me not to: I have a severe case of laryngitis, which has rendered me mute.

Are we to say that I don't have an obligation to be silent? That doesn't seem right at all. It may well be that, since it would make no difference to my deliberation about whether or not to be silent (again, I wouldn't deliberate about that in the first place), it would be pointless to advise or remind me of my obligation, but that doesn't mean it doesn't *exist*. Rather, it is merely that I am in no position to violate it. Worse still, as I just pointed out, if one is not obligated to ϕ , then one is permitted to not- ϕ . So if I have no obligation to be silent, then I am permitted *not* to be silent. So I am permitted to speak up! That is an absurd result. But if we accept the action guidance argument, we are forced to accept it. So we should, I think, reject the action guidance argument.

⁶ If we put aside the act/omission distinction (as I will, since it is not important here) not- ϕ -ing is equivalent to performing another act, ϕ^* . So if you are permitted to not- ϕ if you are unable to not- ϕ , then you are permitted to ϕ^* if you are unable to ϕ^* . Thus we must conclude that inability entails permissibility.

In the next section I'll look at some ways in which action guiders might want to respond to this problem. But before I do, it is worth mentioning a potential amendment to the argument that may seem to avoid the problem. As a referee for this journal has pointed out to me, it isn't only important that we do the right thing, but also that we do so *for the right reasons*. For example, if I do not form an intention to stay silent out of respect for the dead in the *Football* case, then even if, as a matter of fact, I do stay silent, I am not doing the right thing for the right reason; I'm merely doing it because I can't do otherwise. But in order to stay silent out of respect for the dead I must form a contingency plan: to stay silent even if I suddenly gain the ability to speak up. If so, then the would-be obligation to stay silent is guiding for me after all, because it plays a role in guiding my contingency planning. By contrast, one might think, in cases in which I have a would-be obligation to do what I am *unable* to do (rather than what I am unable to *avoid* doing), there is no such guiding role for the would-be obligation to play; if I cannot meet my would-be obligation then, a fortiori, I cannot meet it for the right reasons.

Couldn't action guiders appeal to this asymmetry to argue that a would-be obligation to ϕ can be genuine when one is unable to refrain from ϕ -ing, but not when one is unable to ϕ ? I don't think so. To see why, tweak the *Football* case a bit. Suppose that instead of being mute, I have Tourette's syndrome, and so I am unable to stop myself from shouting out at moments when it would be better to stay silent (we may suppose that the stress brought on by these kinds of moments aggravates my condition.) Even in this case, in which I am unable to stop myself from shouting out, I can nevertheless 'fail better' from a moral point of view by forming an intention not to *deliberately* shout out. And I can do so by forming a contingency plan not to shout out *if* I am able to stay quiet. In this case the impossible-to-satisfy obligation to stay silent is guiding in virtue of the fact that it motivates me to form this contingency plan. But if so, then action guidance considerations don't motivate the idea that there cannot be a moral obligation to ϕ when one cannot ϕ , and so we still don't have a good argument for the oic principle on our hands.⁷

⁷ It has also been suggested to me by a referee that action guiders can avoid a forced march to the absurd conclusion that everything impossible is morally permissible by replacing premise (4) above with (4*): *If you are morally obligated to ϕ , then in so far as deliberation about whether or not to ϕ is rationally permissible, then you can be guided in your (rational) deliberation about whether or not to ϕ by the fact of your obligation to ϕ .* I have to confess that I do not see how an argument can be derived from this premise that is both valid and avoids a forced march to the absurd conclusion.

§III. Responses Considered

The amended argument doesn't work, then. But might action guiders have a good response to the problem I have presented for the (4)-(6) version of the argument? Let's go through their options.

§3.1. Option 1: Argue That It Is True That You Are Permitted To ϕ

One possibility might be to respond by embracing the conclusion that one is not obligated to do what one cannot avoid doing and putting forward an error-theoretic explanation for why it seems wrong. The hypothesis of the error theory is this: strictly speaking it is true that when I cannot avoid ϕ -ing I don't have an obligation to ϕ , but when we *say so* we misleadingly implicate (in the pragmatic, Gricean, sense) a falsehood: that were I *able* to not- ϕ , I would be permitted to do so. It is this misleading implication, the response goes, that explains why the claim sounds absurd even though it is in fact true. So for example, it is strictly speaking true that I don't have an obligation to stay quiet at the football match, but when we say this we generate an implicature to the effect that if I were able to speak up I would be permitted to do so, which is false. And that's why the claim sounds wrong.

I do not think that this response is a good one. Whilst it may seem *prima facie* appealing when it comes to the claim that I don't have an obligation to stay quiet during the requested silence, when we conjoin this claim with the duality principle $\neg O \phi \leftrightarrow P \neg \phi$ it follows that I am permitted *not* to stay quiet. Can the error-theoretic approach explain why this sounds bad? Surely not. It is not remotely plausible to suggest that strictly speaking I *am* permitted to speak up, but that the apparent absurdity of this claim can be explained by the fact of a false implication that I would be permitted to speak up were I able to do so. If someone were to say 'He is permitted to speak up, but I don't mean to suggest that he would still be permitted to do so were he actually able to' the appropriate response wouldn't be assent, but rather to ask: in virtue of what is he permitted to speak up? The answer, of course, must be: in virtue of his inability to do so, since everyone who *hasn't* been rendered mute is obligated to stay quiet. So now we are back to the claim that agents are morally permitted to do everything they are unable to do. But that remains an utterly bizarre claim. ϕ -ing does not become permissible simply in virtue of being impossible.⁸

⁸ It might be tempting to think that statements of the form 'S is permitted to ϕ ' don't pragmatically implicate that S would be permitted to ϕ if he was able to, but rather *express the proposition* that 'S is permitted to ϕ if he/she is able to'. But that can't be right. If it was, then to assert 'He is permitted to speak up, but I don't mean to suggest that he would still be permitted to do so were he actually able to'

For this reason action guiders who want to respond to the objection will have to take a different tack. As I see it there are two options. The first is to argue that in addition to it being false that I am obligated to stay quiet, it is also false that I am permitted not to stay quiet (i.e. to speak up). The second is to argue that when one is unable to ϕ , statements of the form 'S is permitted to ϕ ' *lack a truth-value*. Let's take these in turn.

§3.2. Option 2: Argue That It Is False That You Are Permitted To ϕ

Perhaps action guiders can argue that it is both false that I am obligated to stay quiet and false that I am permitted to speak up? The problem with this proposal is (rather obviously) that it leads to a contradiction when conjoined with the duality principle. To see this, first we take the left-to-right direction of that principle:

$$7. \neg O \phi \rightarrow P \neg \phi$$

By hypothesis I am not obligated to stay quiet, so we have:

$$8. \neg O \phi$$

By modus ponens it follows that:

$$9. P \neg \phi$$

But by hypothesis I am not permitted to speak up (i.e. to not stay quiet), so:

$$10. \neg P \neg \phi$$

And (9) and (10) contradict one another. The upshot is that in order to make this response work one must motivate rejecting the principle $\neg O \phi \leftrightarrow P \neg \phi$. That, I think, will be very hard to do. The principle is among the most basic and most plausible in deontic logic. What reason could one have for rejecting it in this context, other than to save the action guidance argument? I can think of none. But rejecting it simply in order to save the argument would be *ad hoc* and unmotivated.⁹

would be tantamount to asserting 'He is permitted to speak up if he is able to, but I don't mean to suggest that he is permitted to speak up were he able to', which is clearly self-contradictory. But the assertion 'He is permitted to speak up, but I don't mean to suggest that he would still be permitted to do so were he actually able to' doesn't sound self-contradictory, it sounds fine. So the suggestion that 'S is permitted to ϕ ' expresses the proposition 'S is permitted to ϕ if he/she is able to' is not plausible.
⁹ To be clear, I don't mean to suggest that the duality principle should be treated as indefeasibly immune to the possibility of revision. As Sayre-McCord (1986) has pointed out, *no* principle of deontic

§3.3. Option 3: Argue That 'You Are Permitted To ϕ ' Has No Truth Value

Given this problem with the idea that it is *false* that I am permitted to speak up, and given the obvious falsity of the claim that it is *true* that I am permitted to speak up, action guiders are very much under pressure to adopt the second line of argument: that when one is unable to ϕ the proposition that one is permitted to ϕ *lacks a truth value* (as does the proposition that one is obligated to not- ϕ). Here they might take inspiration from a natural way of thinking about the normative status (or lack thereof) of actions performed by non-human animals. Because he isn't a moral agent, it isn't true that my friend's dog, Floyd, has a moral obligation not to bite me. But we shouldn't infer from this that he is morally permitted to bite me. And we shouldn't conclude from the fact that it isn't true that he's obligated not to bite me and isn't true that he's permitted to bite me that the duality principle is false. Rather, we should say that the principle has no application in this case because the propositions 'Floyd is not obligated not to bite me' and 'Floyd is permitted to bite me' lack a truth value: the actions of dogs aren't the *kind of thing* that can be obligatory or not obligatory, permitted or not permitted, and so on. Only the actions of moral agents can be. If it can be argued that 'S is permitted to ϕ ' lacks a truth value when S cannot ϕ in the same way that 'my friend's dog is permitted to bite me' does, then action guiders can retain the duality principle but argue that it has no application in the kinds of cases I have presented as a problem for the action guidance argument. They would thereby avoid having to say that one is permitted to do everything that one cannot do.

logic should be. However, it is one thing to think that the principle is potentially revisable, if the price is worth paying, and quite another to think that the price is in fact worth paying. It is, I think, doubtful that a good case can be made for rejecting the principle in this context. At any rate, it is incumbent on action guiders to make the case if there is one to be made. On this issue, a referee has pointed out to me that moral error theorists such as Olson (2014) and Pigden (2007), who deny that there are any moral properties, reject the duality principle, and that they do not do so in order to save the oic principle. Rather, they reject it in order to make sense of the idea that statements of the form 'S is morally obligated to ϕ /not- ϕ ' and 'S is morally permitted to ϕ /not- ϕ ' are universally false, which would be an incoherent view if the duality principle were sound. Can this line of argument help rescue the action guidance argument? I don't think so. For one thing, it is far from clear that a view according to which moral statements are universally false is in fact coherent. But even putting that aside, moral error theorists will (or at least, should) accept a restricted version of the duality principle according to which, *if* ϕ -ing has moral properties, then ϕ -ing is not obligatory iff not ϕ -ing is permitted (i.e. $(M\phi) \rightarrow (\neg O\phi \leftrightarrow P\neg\phi)$.) Of course, moral error theorists will think that the antecedent of this conditional is always false, but that in itself provides them with no reason to reject the conditional itself. The important point here is that, as we will see, everything I want to do in this article can be done with this restricted principle. This is because, as I'll show in §3.3, action guidance considerations provide us with no reason to reject the idea that ϕ -ing can be morally permissible even when one cannot ϕ (or cannot but ϕ .) *A fortiori*, action guidance considerations give us no reason to think that impossible actions lack moral properties.

Now, taking this line obviously requires showing (at the very least) that action guidance considerations not only motivate the claim that obligation entails ability, but also the claim that *permission* entails ability. For if it is not the case that 'S is permitted to ϕ ' entails 'S can ϕ ', then 'S is permitted to ϕ and S cannot ϕ ' is possibly true, which is precisely what this line of argument denies. So, do action guidance considerations motivate the claim that permission entails ability?¹⁰ Herein lies the problem: they do not. The thought driving the action guidance argument is that no rational agent would deliberate about whether or not to ϕ if they could not ϕ , as such deliberation would be pointless. Since, the argument goes, it is an essential feature of an obligation to ϕ that it guides deliberation about whether or not to ϕ , it follows that a would-be obligation to ϕ is genuine only if one can ϕ . So the question is: should we take the fact that no rational agent would deliberate about whether or not to ϕ if they could not ϕ to also show that a *permission* to ϕ is genuine only if one can ϕ ? The answer is that we should not.

Here's why. There is a vast (possibly infinite) number of things that I am morally permitted to do and can do in the next five minutes, but have no *reason* to do. For example, I am both morally permitted and able to spend the next five minutes counting the number of flies in my office whilst hopping on one leg and singing The Human League's 'Don't You Want Me' (call this the 'counting flies exercise'.) But I have no reason whatsoever to do so. For many of the things that I have no reason to do, I also have no reason to deliberate about whether or not to do them. This is true of the counting flies exercise. It would be a complete waste of time for me to deliberate about whether or not to do that when I have more important things to be thinking about. Given that rationality requires that one not do what one knows one has no reason to do, rationality requires that I don't deliberate about whether or not to perform the counting flies exercise. This is especially true if I have a pressing practical need to think something else right now. So if I am

¹⁰ Since our aim is to see if the action guidance argument can be rescued from the problem I have presented for it, it is crucial that the motivation for the permission-entails-ability principle comes from *action guidance* considerations, rather than something else. Given that 'S is obligated to ϕ ' plausibly entails 'S is permitted to ϕ ', if our reason for thinking that 'S is permitted to ϕ ' lacks a truth value when S cannot ϕ comes from something other than action guidance considerations, then what we will have on our hands won't be a shoring up of the action guidance argument, but rather a wholesale replacement of it with something else, for the considerations that motivated the idea that permission-entails-ability would also motivate the oic principle all by themselves. The action guidance argument would thereby be rendered superfluous; a spinning wheel without any dialectical grip of its own. But as I said earlier, my aim here isn't to evaluate every possible argument for the oic principle; it is only to evaluate the action guidance argument in particular. One might, for example, argue that permissibility-entails-ability on broadly Kantian grounds. But given that such a Kantian argument, if sound, would motivate the oic principle all by itself, it cannot be used to shore up the action guidance argument; instead it would simply render that argument redundant. It wouldn't be *action guidance* considerations motivating the oic principle; it would be the Kantian argument. Thanks to two anonymous referees for pointing out the need for clarification on this point.

rational, I won't deliberate about whether or not to perform it. But it remains true that I am morally permitted to perform it. After all, we surely don't want to deny that there are actions that I am morally permitted to perform and have no reason to perform. What this shows is that one can be morally permitted to ϕ despite the fact that it would be irrational for one to deliberate about whether or not to ϕ , because one has no *reason* to ϕ or to deliberate about ϕ -ing. But in that case we cannot object to the idea that one can be morally permitted to do what one *cannot* do on the grounds that it would be irrational to deliberate about whether or not to do it. For if that is the motivation for concluding that 'S is permitted to ϕ ' entails 'S is able to ϕ ', then, by parity of reasoning, we would also have to conclude that 'S is permitted to ϕ ' entails 'S has a reason to ϕ ', because in so far as it is irrational to deliberate about whether or not to do something one is unable to do, it is equally irrational to deliberate about whether or not to do something when one knows that one has no reason to engage in such deliberation. But 'S is permitted to ϕ clearly does not entail 'S has a reason to ϕ '. Hence, there is no reason to think that permissions are action guiding in a way that motivates the view that permissibility entails ability.¹¹

But if action guidance considerations don't motivate the idea that permissibility entails ability, then the 'no truth value' view doesn't even get off the ground, and we are still left with the idea that there is a truth value to the proposition that I am permitted to speak up. In the *Football* case the truth-value of this proposition is surely 'false'. Since a consistent application of the ideas driving the action guidance argument requires us to say that it is *true*, the argument must be rejected.

§3.4. *Where Does The Action Guidance Argument Go Wrong?*

We've looked at several ways in which action guiders might try to avoid committing themselves to the conclusion that ϕ -ing is permissible whenever it is impossible. None of them works. I conclude that the action guidance argument fails: the only plausible interpretation of it leads to absurd consequences. Where does it go wrong then? The answer must surely be with the claim that an obligation to ϕ *necessarily* plays the role of guiding one's deliberation about whether or not

¹¹ Action guiders might want to respond to this argument by saying that a permission to ϕ is guiding when I have no reason to ϕ because I *could* have had a reason to ϕ , and in the counterfactual world in which I do the fact that I am permitted to do so can serve to guide my deliberation. But this won't help. Even if I lack the ability to ϕ , I *could* have had the ability to ϕ (except if ϕ -ing is metaphysically impossible.) So if we are going to allow guidance in counterfactual worlds into our understanding of what it is for a permission to be guiding, then we can have no objection to the idea that a permission to ϕ is guiding even when one cannot ϕ , because it guides deliberation in counterfactual worlds in which one *can* ϕ . So the idea that action guidance considerations motivate the view that permissibility entails ability is still unmotivated.

to ϕ , for it was this claim that lead us to the absurd conclusion that ϕ -ing is permissible whenever it is impossible. The claim reads too much into the guiding role of obligations. No doubt there are *some* contexts in which an obligation to ϕ can guide one's deliberation about whether or not to ϕ . But this need not always be the case. There are also contexts in which one is obligated to ϕ even though the fact of this obligation cannot make any difference to one's deliberation. The *Football* case shows as much.

§IV. A Proposal

In light of this, and of the failure of the action guidance argument, should we give up on the whole idea of morality being action guiding? I think that would be too hasty. There *is* a particular, broader, sense in which we might reasonably expect moral obligations to be action guiding, but it doesn't motivate the view that one can be obligated to ϕ only if one can ϕ .

To see what I have in mind, notice that we don't just face one-off moral obligations of the kind the discussion has focused on so far; we are also subject to conditional obligations that take the form of rules.¹² An obvious example is something like this:

Promise Rule: If you make a promise, you are obligated to keep it

Promise Rule is almost certainly too simplistic, but it is likely that something in the vicinity is true. We need not worry about the details here, as I only want to use it for illustrative purposes. To begin with notice that if it is true, then the oic principle is false. It is possible to make a promise that you cannot keep, so if, as *Promise Rule* states, you are obligated to keep your promises, then it is possible to be obligated to do something you cannot do. Now, it may be that *Promise Rule* is not deliberatively guiding in contexts in which you have promised to ϕ and you cannot ϕ . But there is a perfectly good sense in which it can guide your deliberation despite this, since it is not only in these contexts that you can use the rule to guide your deliberation about how to act. If you know that *Promise Rule* is in force, and so know that if you make a promise, you are obligated to keep it, then you can use this knowledge to guide your deliberation about whether or not to make promises you know (or think it is likely) you won't be able to keep. If you are in the USA, for example, you can use your knowledge of *Promise Rule* to guide your deliberation about whether or not to promise to meet a friend for lunch

¹² At least, assuming that moral particularism is false.

in London in an hour. You can also use it to guide your deliberation about whether or not to take steps to ensure that you are in a position to keep a promise that you have already made. If you have promised to help a friend move house in the USA tomorrow, for example, then you can use your knowledge of *Promise Rule* to guide your deliberation about whether or not to catch a flight to London today.

These are not contexts in which one has promised to ϕ and cannot ϕ . And as I said, perhaps *Promise Rule* is not deliberatively guiding in such contexts. But, as the above points show, that doesn't mean there are *no* contexts in which it can provide deliberative guidance. Now, it may be that, of necessity, a moral rule is deliberatively guiding in *some* contexts. If so, then a moral rule that could not provide deliberative guidance in any context would not be a genuine moral rule at all. I think that is quite plausible, though I won't defend the claim here. The important point is that if it is true, then morality is, in some sense, necessarily action guiding. But of course this provides us with no reason to think that the oic principle is true.

§V. Conclusion

Morality may or may not be action guiding of necessity. But there are a number of different things one might mean by this claim. The action guidance argument, in what seems to me to be its most (and indeed, only) plausible form, requires an interpretation of it that we must reject, for it leads to the absurd conclusion that everything impossible is morally permissible. Of course, it may be that those who endorse the action guidance argument will reject my characterisation of it. But if so, it is incumbent on them to explain how else it should be understood and why we should accept it.¹³

§VI. References

Vuko Andric, 'Objective Consequentialism And The Rationales Of 'Ought' Implies 'Can'' *Ratio* 29 (2015) pp. 1-16

David Copp, 'Defending the Principle of Alternate Possibilities: Blameworthiness and Moral Responsibility' *Nous* 31 (1997) pp. 441-456

David Copp, "'Ought' Implies 'Can', Blameworthiness, and the Principle of

¹³ Thanks to Daniel Greco, Torfinn Huvenes, Herman Cappelen, Juhani Yli-Vakkuri, Peter Fritz, an audience at University College Dublin, and two anonymous referees for *Utilitas* for helpful comments on earlier drafts of this article.

Alternate Possibilities.' *Moral Responsibility and Alternative Possibilities: Essays on the Importance of Alternative Possibilities*, ed. M McKenna and D Widerker (Burlington, VT, 2003) pp. 265–99.

Julia Driver, 'Promises, Obligations, and Abilities' *Philosophical Studies* 44 (1983) pp. 221-223

Christopher Gowans, *Moral Dilemmas* (New York, 1987)

Peter Graham, 'Ought' and Ability" *Philosophical Review* 120 (2011) pp. 337-382

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 J. Paul (Cambridge: 1992)

Ishtiyaque Haji, *Moral Appraisability: Puzzles, Proposals, and Perplexities* (New York: 1998)

Ishtiyaque Haji, "Moral Anchors and Control" *Canadian Journal of Philosophy* 29 (1999) pp. 175-203

Ishtiyaque Haji, *Deontic Morality and Control* (New York: 2002)

R. M. Hare, *Reason And Freedom* (Oxford: 1963)

Christopher Jay, 'Impossible Obligations Are Not Necessarily Deliberatively Pointless' *Proceedings of The Aristotelian Society* 113 (2012) pp. 381-389

Wayne Martin, 'Ought But Cannot' *Proceedings of The Aristotelian Society* 109 (2009). pp. 103-128

H. E. Mason, "Moral Dilemmas and Moral Theory" (Oxford: 1996)

Elinor Mason, "Consequentialism and the 'Ought Implies Can' Principle" *American Philosophical Quarterly* 40 (2003) pp. 319-331

Jonas Olson, *Moral Error Theory: History, Critique, Defence* (Oxford: 2014)

Charles Pigden, "Nihilism, Nietzsche and the Doppelganger Problem" *Ethical Theory and Moral Practice* 10 (2007). pp. 441-456

Geoffrey Sayre-McCord, "Deontic Logic and The Priority of Moral Theory" *Nous* 20 (1986). pp. 179-197

Walter Sinnott-Armstrong, 'Ought Conversationally Implies 'Can'' *The Philosophical Review* 93 (1984). pp. 249-261

Robert Stern, 'Does 'Ought' Imply 'Can'? And Did Kant Think It Does?' *Utilitas* 16 (2004). pp. 42-61

Peter van Inwagen, *An Essay on Free Will*. (Oxford: 1983)

David Widerker, 'Frankfurt on 'Ought Implies Can' and Alternative Possibilities' *Analysis* 51 (1991) pp. 222-224

Bernard Williams, 'Internal Reasons and the Obscurity of Blame' *Making Sense of Humanity* (Cambridge: 1989)

nickhowellhughes@gmail.com