

Forthcoming in *The Philosophical Quarterly*

Daniel Immerman

Does Ought Imply Ought Ought?¹

Abstract. Knows-knows principles in epistemology (KK principles for short) say that if you know some proposition, then you are in a position to know that you know it. This paper examines the viability of analogous principles in ethics, which I call ought-ought principles (OO principles for short). Several epistemologists have recently offered new defenses of KK principles and of other related principles and there has recently been an increased interest in examining analogies between ethics and epistemology and so it seem natural to examine whether defenses of KK and related principles carry over to OO principles. In this paper I introduce two OO principles and I show how some arguments in favor of KK carry over to them. Then I show how these OO principles can be used to shed light on a much-discussed case in ethics, that of Professor Procrastinate.

Knows-knows principles in epistemology (KK principles for short) say that if you know some proposition, then you are in a position to know that you know it.² This paper will examine the viability of analogous principles in

¹Thanks for helpful comments to Emily Dawson and two anonymous referees.

²For discussion of KK principles, see e.g. [Castañeda, 1970], [Conn, 2001], [Das and Salow, Forthcoming], [Feldman, 1981], [Ginet, 1970], [Greco, 2014], [Greco, 2015] and [Hemp, 2014].

ethics, which I will call ought-ought principles (OO principles for short).³

Here are three reasons to be interested in the project of examining the viability of OO principles. First, there has been a large amount of work done on the structural relationships among first-order moral obligations, for instance, on whether it is possible to have conflicting moral obligations.⁴ But there has been comparatively less work done on the relationship between first-order and higher-order moral obligations, so this area is ripe for investigation.

Second, several epistemologists have recently offered new defenses of KK principles and of other related principles.⁵ Meanwhile, there has recently been an increased interest in examining analogies between ethics and epistemology. Given these developments, it seem natural to examine whether defenses of KK and related principles carry over to OO principles.

Third, the issue of whether there are viable OO principles is relevant to a longstanding debate in ethics. The debate, which is sometimes known as the actualism/possibilism debate, involves cases in which characters are trying

³Such principles are rarely discussed in the ethics literature. That said, some discussion of them has appeared in works on deontic logic, such as [von Wright, 1983]. It is worth noting that deontic logic provides a potentially interesting source of arguments for ought-ought principles. There is a standard system of modal logic, S4, in which $\Box p \rightarrow \Box \Box p$. And in deontic logic, obligation operators are necessity operators and thus, there is a corresponding system of deontic logic on which $Op \rightarrow OOp$ [von Wright, 1983, 126]. In other words, there is a corresponding system in which the OO principle is true. So there is some work to be done looking at the arguments for S4 and whether they carry over to this corresponding system in deontic logic. While there are interesting issues to be pursued here, they are sufficiently different from the rest of this paper that I shall set them aside. But thanks to an anonymous referee for suggesting I mention them.

⁴This debate seems to have been kicked off by Bernard Williams in [Williams, 1965]; key works include [Conee, 1982] and [Marcus, 1980]. For a good list of sources on the debate, see the bibliography of [McConnell, 2014].

⁵See e.g. [Das and Salow, Forthcoming], [Greco, 2014], and [Greco, 2015].

to decide whether to take on obligations that they can but won't fulfill.⁶ For instance, one such character, Professor Procrastinate, is invited to write a book review; it would be best if he accepts and then writes it, but in fact, if he accepts, he will procrastinate and fail to. As we shall see, the arguments used in favor of OO yield new insight regarding this and other similar cases.

This paper has three sections. In the first, I introduce two OO principles. In the second, I show how some arguments in favor of KK carry over to these OO principles. In the third, I show how these OO principles can be used to shed light on the Professor Procrastinate case.

1 Introducing OO

In this section, I will introduce the two OO principles I'll be focusing on. The OO principles I am discussing use the word "ought" so I will start by clarifying what sense of "ought" I am interested in. I will be using the all-things-considered objective moral sense of "ought". By this I mean the moral ought that takes all relevant information into account, not merely the information that the subject possesses, and is not overridden by some other moral ought. I will use the word "obligation" in this paper interchangeably with the word "ought".

I have chosen to focus on this particular sense of "ought" because I think this sense of "ought" nicely parallels "knows". For instance, just as with my

⁶For some key works, see e.g. [Feldman, 1986], [Goble, 1993], [Goldman, 1976], [Goldman, 1978], [Greenspan, 1978], [Humberstone, 1983], [Jackson and Pargetter, 1986], [Sobel, 1976], [Sobel, 1982], [Thomason, 1981], [Vessel, 2009], [Vorobej, 2000], [Zimmerman, 1996].

sense of “ought”, “knows” is objective – knowing something requires more than merely believing it; it also requires, for instance, that your belief be true. That said, for all I say in this paper, there may well be interesting arguments for OO principles involving subjective oughts.

One issue in stating OO principles – which perhaps explains why they have not been widely discussed – is that it is not immediately clear how to talk about higher-order oughts. It makes sense to say “Anita ought to walk the dog” but it makes less sense to say “Anita ought to ought to walk the dog.”

While it may thus appear that there are no higher-order oughts, in fact, there are a bunch of different ones; ones concerned with taking on obligations, ones concerned with ensuring that one fulfills obligations, ones concerned with acting on one’s obligations in the right way, and so on. For the sake of focus, I will focus on just two OO principles in this paper.

The first one is connected with cases in which we decide whether to take on obligations. For instance, suppose Roger asks Anita if she will take on the obligation to walk the dog later today. Suppose further that she ought to say “yes”. Here, it seems plausible to say that she ought to take on the obligation to walk the dog.

Here then is a first stab at an OO principle:

(OO_{Acquire} – First Stab): If one ought to ϕ then one ought to acquire an obligation to ϕ .

Here is a problem with this principle: sometimes, even if one ought to do something, there may be no opportunity to take on an obligation to do it. So, for instance, suppose that Anita ought to walk the dog. No one else has asked her to walk it or even invited her to. But she's the only one home and the dog needs walking. Here, even though Anita ought to walk the dog, it's not the case that there is some additional thing she ought to do, namely to take on the obligation to walk it. She should just walk it. She doesn't have to hunt someone down and promise to walk it, or in some other way take on the obligation to do so.

One difference between this case and the original one is that in the original case Roger asked Anita if she would walk the dog and in the new case he didn't. In other words, in the original case, she was confronted with the decision of whether to take on the obligation and had to decide one way or the other. So we can modify the original version to yield:

OO_{Acquire}: if one ought to ϕ and one has to decide whether to acquire an obligation to ϕ , one ought to acquire an obligation to ϕ .

Let me add some clarification. My story involved someone promising to do something and thereby acquiring an obligation to do it. But I want to be clear that I am not assuming that promising to do something always results in acquiring an obligation to do it. Sometimes, such as in the Anita case, it will. But in others, it won't – after all, some promises shouldn't be kept.

Likewise, I am not assuming that the only way to acquire an obligation is to promise to do it; there are many other ways to do so.⁷

Also, there is a puzzle involving three claims, each of which I endorse; addressing it will further clarify $OO_{Acquire}$. Here are the three claims: Claim one: when asked if she'll walk the dog, Anita ought to say yes and in so doing she will thereby acquire an obligation to walk it. Claim two: even before saying yes, it was already the case that Anita ought to walk the dog. Claim three: my sense of "ought" is an all-things-considered sense of ought.

It seems hard to square these claims. If one already has an all-things-considered obligation to do something, in what sense can one, at this point, acquire an all-things-considered obligation to do it? Saying that this is possible seems a bit like saying that I can bring it about that the lights are on after the lights are already on.⁸

This presents an interesting puzzle. The first thing to say is that, despite the fact that they can generate a puzzle, the three claims really do all seem true to me and to others I've talked with. So it would be nice to find a solution to this puzzle that retains all three.

I must admit that I find it a little unclear how to do so, but here is a stab. We sometimes individuate obligations via the content of the obligation. In our case, the content is walking the dog. On this way of individuating, Anita has the same obligation – to walk the dog – both before she promises

⁷Thanks to an anonymous referee for suggesting I clarify this.

⁸Thanks to an anonymous referee for pressing this worry.

and after. A second way of individuating an obligation is more fine-grained; it says that obligations with different grounds are different obligations. This will yield the result that there are two different obligations in the Anita case. In particular, before she promised, Anita had an obligation to walk the dog grounded in her responsibility to take care of it. Afterwards, she had an obligation grounded, at least in part, in her promise. If we individuate obligations in this second way, we can see how Anita could both have an obligation to do something and then bring it about that she has an obligation to do this thing. In particular, we are talking about two different obligations, both with the same content, but with different grounds.

A similar thing can be said of the lights being on. Suppose, for example, that it is closing time at the store I work at. Right now, the lights are on. This is thanks to someone else – Lucia – turning them on in the morning. As I am leaving, I have to lock the door, and in doing so, I have to choose between pressing one of two buttons, one of which will keep the lights on at night, and one of which will turn them off. I am supposed to turn the lights off; the owner doesn't want to waste electricity. But I negligently press the wrong button, thus keeping the lights on all night. In the morning, the owner walks in and angrily asks: "Who did this? Whose action brought it about that the lights are on?" It seems appropriate to say that it was my pressing the button to leave the lights on that brought this about. After all, if I had pressed the other button, the lights would now be off. By contrast, it would be somewhat strange to say that it was Lucia who brought it about that the

lights are on.

We have seen one sort of higher-order ought. But there are others. A second sort of higher-order ought concerns cases in which one is trying to decide whether to ensure that certain obligations are fulfilled. For example, suppose that Anita has the moral obligation to walk her dog this afternoon. In such a case, it seems somewhat plausible to say that she also ought to ensure that her obligation is fulfilled. So, for instance, if she is likely to forget to walk the dog, then she ought to set an alarm to remind herself.

This yields a second OO principle, viz.

OO_{Ensure}: if one has an obligation to ϕ , one ought to ensure that one's obligation is fulfilled.

2 Arguments for KK and their analogues for OO

In this section, I will look at several key arguments in defense of KK principles and similar principles. My goal is to argue that these arguments carry over to my OO principles. The arguments I will be discussing are quite contentious and have been much discussed. In my discussion here, I will not aim to fully defend the premises of these arguments. Rather, I shall work to argue that if the arguments are plausible in the case of KK, then the analogous arguments in favor of OO are plausible.

2.1 KK, Internalism, and OO_{Ensure}

One significant argument in favor of KK comes from “internalism about knowledge”. I shall first introduce the argument and then show how an analog of it can be used to defend OO_{Ensure} .

By “internalism about knowledge”, as I shall be using the term, I mean that a major component of knowledge – I will sometimes refer to this component as “justification” – supervenes on “internal properties”. There are several things that people in the literature mean by “internal properties”. On one understanding, “access internalism”, internal properties are those that are internally accessible.⁹ There are different understandings of “internally accessible” properties; this can be understood as those that are accessible by reflection, or those that are accessible a priori. On another understanding, “mentalism”, internal properties are those that supervene on one’s mental states.¹⁰

Why endorse internalism about knowledge? One argument is that it nicely explains various examples.¹¹ For one, if internal properties differ, then epistemic ones arguably do as well. For instance, suppose you and I are both looking at birds, but I have a visual experience as of seeing a bird and you don’t – you just see what looks like a blue sky. Then I am justified in thinking there’s a bird in front of me and you’re not; I can know that there’s

⁹See e.g. [Audi, 2011], [BonJour, 2010].

¹⁰See e.g. [Conee and Feldman, 2004].

¹¹For this sort of argument, see e.g. [Conee and Feldman, 2004].

a bird in front of me and you can't. Meanwhile, if internal properties stay the same, then key epistemic ones arguably do as well. For instance, suppose I am looking at a bird and you have an identical visual experience but are in fact looking at a stuffed animal. Then, arguably, if I am justified in thinking there's a bird in front of me, you are too. In short, there is evidence from examples that seems to suggest that internalism is true; that justification, a key component of knowledge, supervenes on internal properties.

Another argument in favor of internalism connects to the so-called "deontological conception of justification".¹² This conception – really several tightly related conceptions – says that justification is a matter of fulfilling one's epistemic duty; or relatedly, that whether one is justified in believing something turns on whether one is blameless in believing it; and further that whether one is fulfilling one's epistemic duty and/or whether one is blameless turns on internal properties; after all, one cannot be blamed for something outside of one's awareness or something beyond one's mental states.

Internalism about knowledge can be used to generate an argument for KK.¹³ A key premise of this argument is that we have special access to internal properties. On versions of access internalism this is obvious; internal properties are defined as those we have special access to. But on various versions of mentalism this is also plausible; many think that we have some sort of special access to mental properties.¹⁴ Given that we have special

¹²For this sort of argument, see e.g. [Goldman, 1999]. Note: epistemologists' use of "deontological" here does not perfectly line up with how ethicists typically use it.

¹³See e.g. [Hemp, 2010].

¹⁴Of course, this is controversial; for critical discussion, see e.g. [Alston, 1971], [Heil,

access to these properties, it seems to follow that if we bear them to some proposition, there is no barrier to us knowing that we bear them to that proposition, and thus that some version of KK is true.

To take an example, suppose that I have a visual experience as of a bird in front of me. Because this is an internal state, I have special access to it; I can know that I'm having this visual experience. And thus, not only am I justified in believing that there is a bird in front of me (on the basis of this visual experience), I am also justified in believing that I'm justified in believing that there is a bird in front of me (on the basis of my awareness of this visual experience).

An analog of this argument can be used to defend OO_{Ensure} . This time, I will be talking about “internalism about the object of moral obligations”, understanding it as the view that the objects of our moral obligations are “internal”.¹⁵ Here again there are two senses of “internal”. On the “access internalism” understanding of “internal”, the objects of our obligations are things we have direct control over. So, for example, take the case of walking the dog. An access internalist would say that, strictly speaking, Anita's obligations concern things she has direct control over. There are several options for what counts as direct control; on some understandings, it will be certain sorts of basic actions, on others, it will be certain sorts of intentions. In short, these access internalists say that, strictly speaking, Anita's obligations

¹⁵This should not be confused with the position known as “moral internalism”.

tion to walk the dog is really the obligation to intend to walk the dog, or to perform certain basic actions directed at walking the dog. On the “mental-ism” understanding of “internal”, the objects of our obligations supervene on mental states. For instance, perhaps our obligations are to have certain intentions, or certain sets of belief-desire pairs.

Again, internalism about the object of moral obligations can be defended by both examples and theory. First examples. As before, it seems that internalism can explain our judgments about certain cases. For instance, if internal properties differ, then arguably moral ones do as well. Suppose you and I are both driving cars. Suppose an animal starts to walk in front of each of us. I try to swerve to avoid it and am successful. You try to run it over, but your car, thanks to some faulty mechanism, swerves, and you avoid it as well. Even though the consequences of our actions are the same – we avoid hitting the animal – arguably I fulfilled my moral obligations and you did not. Meanwhile, if internal properties stay the same, then arguably moral ones do as well. Say, for instance, we are both driving and see the animal. This time, we both try to swerve. I manage to, but you don’t, again thanks to some faulty mechanism outside your control. Here, arguably, we behaved equally well; if I fulfilled my obligations, then you did too. The covariance of internal properties and moral obligations suggests that internalism about the object of moral obligations is true.

And again, internalism about the object of moral obligations can also be defended by appeal to theory. Here the idea is that moral justification is

tightly connected with fulfilling one's moral duty; that whether one is justified turns on whether one is blameless; and further that whether one is fulfilling one's duty and/or whether one is blameless turns on internal properties; after all, one cannot be blamed for something that depends on external factors.

Internalism about moral obligation can be used to generate an argument for OO_{Ensure} . The basic idea is that we have special control over internal properties. On versions of access internalism this is obvious; internal properties are defined as those we have special control over. But on versions of mentalism this is also plausible; many think that we have some sort of special control over mental properties. Given that we have special control over these things, it seems to follow that if we should make them a certain way, then there is no barrier to ensuring that our obligation is fulfilled, and thus that OO_{Ensure} is true.

It might be helpful to illustrate this argument with a concrete example. In fact, there are multiple ways to do so because there are multiple ways to spell out what it means to have special control over internal properties. I gave one concrete example while introducing OO_{Ensure} . Here is a second, which is somewhat different. Suppose that it is indeed the case that Anita ought to walk the dog. Suppose further that the objects of our obligations are internal in the following sense: our obligations are obligations to have certain intentions. Then if it is indeed the case that Anita ought to walk the dog, it follows that there is no barrier to her forming the intention to walk the dog – after all, it is an internal state – and thus that she ought to ensure

that she does form this intention.

One might worry that what I just said is in tension with my claim that I am talking about an objective sense of “ought”. After all, in saying that the moral ought was objective, I said that it didn’t depend merely on information the subject possessed. But now I’m giving an argument that turns on what I’m call internalism about moral obligation. And this says that the objects of our moral obligation are internal. Isn’t it inconsistent to hold that whether we ought to do something depends on external matters, where the thing we ought to do is an internal matter?¹⁶

My response: I think it is possible to consistently hold both. *Whether* we ought to do something is one matter, *what* we ought to do a different matter. Here’s an example: consider the claim that you ought to believe true things and not believe false ones. Whether something is true or false is an objective matter; it doesn’t merely depend on the information one has. And thus whether one ought to believe something – assuming this duty exists – is an objective matter. Nonetheless, the object of the duty – forming a belief – is a mental matter.

2.2 KK, Transparency, and $OO_{Acquire}$

Another significant argument in favor of KK and related principles concerns “transparency”. I shall first introduce the argument and then show how an analog of it can be used to defend $OO_{Acquire}$.

¹⁶Thanks to an anonymous referee for pressing this worry.

A number of epistemologists have endorsed an idea called “transparency about belief.” There are several different formulations of this idea, one popular one is that determining whether you believe that p is a matter of comparing reasons for and against p . Once you’ve settled whether p , you’ve thereby settled whether you believe p . For example, if someone asks: “do you believe the cat is on the couch,” a natural way to respond is to check whether the cat is on the couch. Once you have checked and seen that the cat is on the couch, you can respond “yes.” You has thereby determined whether you believe it; no further inquiry is required. As Gareth Evans, who is often credited with the original statement of the idea, puts the point:

If someone asks me, ‘Do you think there is going to be a third world war?’, I must attend, in answering him, to precisely the same outward phenomena as I would attend to if I were answering the question ‘Will there be a third world war?’ I get myself into the position to answer the question whether I believe that p by putting into operation whatever procedure I have for answering the question whether p . [Evans, 1982, 225].¹⁷

In their defense of this point about transparency, transparency theorists like to cite the point that, in deliberating about whether to believe things, only certain sorts of reasons can move us. For example, being told that someone will give me a million dollars if I believe that the moon is made of cheese

¹⁷See also [Byrne, 2005, 82-3], [Fernández, 2003, 355], [Gallois, 1996, 50-1], [Gordon, 1995, 65], [Moran, 2001, 66], and [Shah and Velleman, 2005, 602].

cannot directly move me to believe that the moon is made of cheese. Only facts bearing on whether the moon is made of cheese can so move me.

Transparency about belief seems to lend credence to KK and other similar principles, such as the principle that if you ought to believe something, then you ought to believe that you ought to believe it.¹⁸ If settling whether you believe some proposition seems to reduce to settling whether that proposition is so, then it seems that you ought to believe some proposition just in case that proposition is so. If we substitute in “you ought to believe that p” for the proposition in question, we get that you ought to believe that you ought to believe a proposition just in case you ought to believe the proposition.

Similar thoughts can be extended to defend $OO_{Acquire}$. In particular, we can label the thesis “transparency about intention” that determining whether you intend to ϕ is a matter of comparing reasons for and against ϕ -ing and that once you’ve settled whether to ϕ you’ve thereby settled whether you intend to ϕ . For example if someone asks me: “Do you intend to walk the dog” and I start wondering about whether I have this intention or not, it’s natural for me to start comparing the reasons for and against walking the dog. Once I’ve settled the matter, decided whether to walk the dog or not, I’ve thereby also determined whether I have the intention to walk the dog.

It’s worth noting that the other point cited in favor of transparency about belief – the point about only certain types of reasons being capable of moving

¹⁸I will focus on the argument for the latter; for the argument for the former, see e.g. [Das and Salow, Forthcoming] and [McHugh, 2010].

us – also carries over. In particular, in deliberating about whether to form certain intentions, only certain sorts of reasons can move us. For example, as Gregory Kavka points out, being told that someone will pay me a million dollars if I form the intention to eat something mildly toxic tomorrow cannot directly move me to form the intention [Kavka, 1983]. Only reasons bearing on the object of the intention – in this case, reasons bearing on whether I should eat the thing or not – can so move me.

Transparency about intention seems to lend credence to $OO_{Acquire}$. In particular, it seems that you can only (properly) take on an obligation to do something if you intend to do it. For instance, if you promise to do something but have no intention of doing so, it doesn't seem like a real promise. And thus settling whether to take on an obligation requires forming an intention to do it. But, as noted above, determining whether to form an intention to perform the action seems to reduce to settling whether to perform the action. So it seems to follow that if one ought to ϕ and one has to decide whether to acquire an obligation to ϕ , one ought to acquire an obligation to ϕ .

3 Professor Procrastinate and OO

The Professor Procrastinate case runs as follows:¹⁹

PROFESSOR PROCRASTINATE. Professor Procrastinate re-

¹⁹This case, or ones similar to it, are discussed in e.g. [Baker, 2012, 641], [Cariani, Forthcoming], [Goldman, 1978, 185-6], [Jackson and Pargetter, 1986, 235], [Littlejohn, 2009], [Portmore, 2011, 151], [Portmore, 2013], [Timmerman, 2015, 1512], [van Someren Greve, 2013, 482-3], [Vessel, 28, 166], [Woodard, 2008, 18].

ceives an invitation to review a book. The best thing that can happen is that he says yes, and then writes the review when the book arrives. However ... were Procrastinate to say yes, he would not in fact get around to writing the review. Not because of incapacity or outside interference or anything like that, but because he would keep on putting the task off. Thus, although the best that can happen is for Procrastinate to say yes and then write ... what *would* in fact happen were he to say yes is that he would not write the review. Moreover ... this ... is the worst that can happen. [Jackson and Pargetter, 1986, 235] [Emphasis theirs].

One way of explaining why this case is puzzling is the following, all three of the following principles seem to be true:

(It's Hard To Avoid Obligations): The fact that one will not fulfill an obligation is not enough, by itself, to keep one from having the obligation.

(Don't Bring About Worse Consequences): If doing one thing will lead to worse consequences than doing something else, one ought not do it.

(Take Necessary Means): If one ought to do something, then one ought to take the necessary means to doing that thing.

But these three principles seem to lead to conflicting implications regarding this case. On the one hand, (It's Hard To Avoid Obligations) seems to

yield the conclusion that Procrastinate should write the review; the fact that he won't fulfill this obligation is not enough to keep him from having it. Meanwhile, (Don't Bring About Worse Consequences) seems to yield the conclusion that he shouldn't accept the invitation; his accepting will lead to worse consequences than his turning it down. But these two conclusions are in conflict with (Take Necessary Means); if he ought to write the review, then he ought to take the necessary means to writing it, namely accepting the invitation.

Using our OO principles, we can shed some light on the case. I'll first offer the verdicts these principles offer on the case, then I'll show how they can resolve the conflict seemingly generated by the three principles above.

First, OO_{Ensure} . Recall that OO_{Ensure} runs as follows: if one ought to ϕ , one ought to ensure that one's obligation to ϕ is fulfilled. Applied to the case at hand, OO_{Ensure} says that if Procrastinate ought to review the book, he ought to ensure that his obligation to review the book is fulfilled.

If we combined OO_{Ensure} with a plausible ought-implies-can principle, we can yield an interesting and plausible verdict on the Procrastinate case. The ought-implies-can principle in question says the following: if, at some time, one ought to do something, then, at that time, one can do it.

This yields a question, is there anything Procrastinate can do now to ensure that he completes the review? The case does not tell us what the answer to this question is. On one reading of the case is that there is nothing Procrastinate can do right now to ensure that he will complete the review.

Even if he has good intentions now, in the future he will fail to act on them, failing to write the review. So, if we combine ought implies can with OO_{Ensure} , we get the conclusion that it is not the case that he ought write the review.

The other reading of the case is that Procrastinate can do something right now to ensure that he writes the review. For example, maybe he can do something to ensure that he doesn't procrastinate, like making a promise that if he fails to write the review, he'll donate a large amount of money to some charity he loathes. In that case, OO_{Ensure} will allow that he ought to accept the invitation, but say that he also ought to ensure that he writes the review.

Some have offered similar verdicts regarding the Procrastinate case, and thus it's worth identifying what my paper adds to the discussion. First: my paper focuses on an objective, all-things-considered moral ought, and thus differs from papers focused on another sort of ought, such as a Brian Hedden's "Options and the Subjective Ought", which focuses on a subjective ought [Hedden, 2012]. Second: my arguments for OO_{Ensure} mostly turn on meta-ethical concerns and are thus – to a significant degree – neutral when it comes normative ethical theory; they are, so far as I can tell, compatible with consequentialism, deontology, and virtue ethics. Thus they contrast with those, like Jacob Ross in "Actualism, Possibilism, and Beyond" [Ross, 2012] and Douglas Portmore in a number of works, e.g. [Portmore, 2011] and [Portmore, 2013], who offer similar verdicts but tie their verdicts to

consequentialism. Thirdly, my paper offers a distinctive approach when it comes to the defense of OO_{Ensure} and thus, even insofar as its verdict is similar to others' verdicts, its reasons for that verdict differ.

Next, let us see how $OO_{Acquire}$ handles the Procrastinate case. Recall that this principle runs as follows: if one ought to ϕ and one has to decide whether to acquire an obligation to ϕ , one ought to acquire an obligation to ϕ . Recall further that the defense of $OO_{Acquire}$ turned on the idea that deliberating whether to take on an obligation to perform an action naturally turns into deliberating about the action itself.

This suggests that in deliberating about whether to take on the obligation to write the review, Procrastinate should deliberate about whether to write it, and should only accept the invitation if he is really committed to writing the review.

Again, this allows us two readings of the case. On one reading, there is no way for him to fully commit to write the review. Even if he agrees to write the review and says he will write it, in fact, he won't be fully committed and will back out. On this reading of the case, he shouldn't accept the invitation, given that he cannot fully commit to write.

On the other reading, it is possible for Procrastinate to fully commit to writing the review. This isn't to say that any old way of accepting the invitation to write the review would amount to a full commitment to write it. Rather, it is to say that certain ways of accepting the invitation do involve a full commitment. And if he does fully commit, then he will write the

review. In short, on this resolution of the case, those who procrastinate do so because they are never fully committed in the first place. In light of this, Procrastinate should only take on the obligation to write the review – and only has the obligation in the first place – if he can be fully committed to writing it.

Now that we have seen the ways in which our OO principles respond to the case, we can turn back to see how they address the three seemingly conflicting principles. First, there was (It's Hard To Avoid Obligations): The fact that one will not fulfill an obligation is not enough, by itself, to keep one from having the obligation. This principle seemingly generated the consequence that Procrastinate ought to write the review; the fact that he won't write it is not, by itself, enough to keep him from having this obligation.

Using our OO principles, we can grant the principle, while denying that it has the implication that Procrastinate ought to write the review. Each of the two OO principles allowed that, on certain readings of the case, it's not the case that Procrastinate ought write the review. But in neither case was it his failure to write it that explained why he lacked the obligation to write it. Rather, OO_{Ensure} only said he lacked the obligation to write it if there was no way he could ensure that he would write it. And $OO_{Acquire}$ said that he lacked the obligation if there was no way for him to be fully committed to writing it.

Next, there was (Take Necessary Means): If one ought to do something, then one ought to take the necessary means to doing that thing. Both OO

principles were consistent with this principle.

Finally, there was (Don't Bring About Worse Consequences): If doing one thing will lead to worse conclusions than doing something else, one ought not do it.

In light of our discussion of the OO principles, it seems that this principle is not quite correct. In particular, there are possible cases in which the reason that doing the one thing will lead to a worse conclusion is that one is doing it the wrong way. In such a case, it may still be the case that one ought to do the thing in question; one just ought to do it in a better way. This is most apparent in cases in which doing the thing in question in a better way is fully in one's control. For example, suppose that I am at dinner and am trying to decide whether to ask someone to pass the peas or just reach across the table and take them. The very best option is for me to ask for them politely. But if I do ask for them, I'll ask for them rudely, which is even worse than reaching across the table and taking them. Just to be clear, I am fully capable of being polite; if I want to be polite, I can. It's just that if I ask for the peas, I won't be polite. In this case, Take Necessary Means suggests that I ought not ask for the peas. But this is (arguably) implausible; I should ask for them, but in a polite manner.

So it appears that this principle was false, and a better principle is (Take Necessary Means*): if the best way of doing one thing will lead to worse consequences than the best way of doing something else, one ought not do it.

This principle is consistent with the peas example; it does not say that I ought not ask for peas, because the best way of asking for peas – asking politely – does not lead to worse consequences than the best way of reaching across the table.

The initial problem with (Take Necessary Means) was that it seemed to yield the conclusion that Procrastinate should not accept the invitation, seeing as his accepting it would lead to his procrastinating and failing to write the review. But (Take Necessary Means*) does not immediately tell us whether he should accept. Rather, whether he should accept turns on what the best way of his accepting would be. As our OO principles revealed, perhaps there are ways of accepting that would lead to him writing the review. In particular, maybe there are ways of accepting that will ensure that he will write, or that will involve a full commitment to his writing. In such a case, he should accept, and if not, not.

In short, accepting either of the OO principles can help us resolve the puzzle of the Professor Procrastinate case. This provides a second reason, apart from the arguments I gave for them earlier on, for accepting these OO principles.

References

[Alston, 1971] Alston, W. P. (1971). Varieties of privileged access. *American Philosophical Quarterly*, 8(3):223–41.

- [Audi, 2011] Audi, R. (2011). *Epistemology: A Contemporary Introduction to the Theory of Knowledge*. Routledge, third edition.
- [Baker, 2012] Baker, D. (2012). Knowing yourself – and giving up on your own agency in the process. *Australasian Journal of Philosophy*, 90(4):641–56.
- [BonJour, 2010] BonJour, L. (2010). Externalism/internalism. In Dancy, J., Sosa, E., and Steup, M., editors, *A companion to epistemology*, pages 364–369. Blackwell, 2nd edition.
- [Byrne, 2005] Byrne, A. (2005). Introspection. *Philosophical Topics*, 33(1):79–104.
- [Cariani, Forthcoming] Cariani, F. (Forthcoming). Consequence and contrast in deontic semantics. *The Journal of Philosophy*.
- [Castañeda, 1970] Castañeda, H.-N. (1970). On knowing (or believing) that one knows (or believes). *Synthese*, 21(2):187–203.
- [Conee, 1982] Conee, E. (1982). Against moral dilemmas. *The Philosophical Review*, 91(1):87–97.
- [Conee and Feldman, 2004] Conee, E. and Feldman, R. (2004). Evidentialism. In *Evidentialism: Essays in Epistemology*. Oxford University Press.
- [Conn, 2001] Conn, C. (2001). Chisholm, internalism, and knowing that one knows. *American Philosophical Quarterly*, 38(4):333–347.

- [Das and Salow, Forthcoming] Das, N. and Salow, B. (Forthcoming). Transparency and the k k principle. *Noûs*.
- [Evans, 1982] Evans, G. (1982). *The varieties of reference*. Oxford University Press.
- [Feldman, 1986] Feldman, F. (1986). *Doing the Best We Can*. D. Reidel.
- [Feldman, 1981] Feldman, R. (1981). Fallibilism and knowing that one knows. *The Philosophical Review*, 90(2):266–282.
- [Fernández, 2003] Fernández, J. (2003). Privileged access naturalized. *The Philosophical Quarterly*, 53(212):352–72.
- [Gallois, 1996] Gallois, A. (1996). *The world without, the mind within*. Cambridge University Press.
- [Ginet, 1970] Ginet, C. (1970). What must be added to knowing to obtain knowing that one knows. *Synthese*, 21(2):163–186.
- [Goble, 1993] Goble, L. (1993). The logic of obligation, ‘better’ and ‘worse’. *Philosophical Studies*, 70(2):133–63.
- [Goldman, 1999] Goldman, A. (1999). Internalism exposed. *The Journal of Philosophy*, 96(6):271–293.
- [Goldman, 1976] Goldman, H. S. (1976). Dated rightness and moral imperfection. *The Philosophical Review*, 85(4):449–87.

- [Goldman, 1978] Goldman, H. S. (1978). Doing the best one can. In Goldman, A. I. and Kim, J., editors, *Values and Morals*. D. Reidel.
- [Gordon, 1995] Gordon, R. M. (1995). Simulation without introspection or inference from me to you. In Davies, M. and Stones, T., editors, *Mental simulation*. Wiley-Blackwell.
- [Greco, 2014] Greco, D. (2014). Could kk be ok? *The Journal of Philosophy*, CXI(4):169–197.
- [Greco, 2015] Greco, D. (2015). Iteration and fragmentation. *Philosophy and Phenomenological Research*, 91(3):656–73.
- [Greenspan, 1978] Greenspan, P. S. (1978). Oughts and determinism: A response to goldman. *The Philosophical Review*, LXXXVII(1):77–83.
- [Hedden, 2012] Hedden, B. (2012). Options and the subjective ought. *Philosophical Studies*, 158(2):343–60.
- [Heil, 1988] Heil, J. (1988). Privileged access. *Mind*, 97(386):238–51.
- [Hemp, 2010] Hemp, D. (2010). Knowledge and conclusive evidence. In Campbell, J. K., O’Rourke, M., and Silverstein, H. S., editors, *Knowledge and skepticism*, pages 27–43. MIT Press.
- [Hemp, 2014] Hemp, D. (2014). The kk (knowing that one knows). *The Internet Encyclopedia of Philosophy*.

- [Humberstone, 1983] Humberstone, I. L. (1983). The background of circumstances. *Pacific Philosophical Quarterly*, 64:19–34.
- [Jackson and Pargetter, 1986] Jackson, F. and Pargetter, R. (1986). Oughts, options, and actualism. *The Philosophical Review*, 95(2):233–55.
- [Kavka, 1983] Kavka, G. S. (1983). The toxin puzzle. *Analysis*, 43(1):33–36.
- [Littlejohn, 2009] Littlejohn, C. (2009). Critical notice of michael zimmerman’s *Living With Uncertainty*. *Philosophical Books*, 50(4):235–47.
- [Marcus, 1980] Marcus, R. B. (1980). Moral dilemmas and consistency. *The Journal of Philosophy*, 77(3):121–36.
- [McConnell, 2014] McConnell, T. (2014). Moral dilemmas. *Stanford Encyclopedia of Philosophy*.
- [McHugh, 2010] McHugh, C. (2010). Self-knowledge and the kk principle. *Synthese*, 173:231–257.
- [Moran, 2001] Moran, R. (2001). *Authority and estrangement*. Princeton University Press.
- [Portmore, 2011] Portmore, D. W. (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford University Press.
- [Portmore, 2013] Portmore, D. W. (2013). Perform your best option. *The Journal of Philosophy*, 110(8):436–59.

- [Ross, 2012] Ross, J. (2012). Actualism, possibilism, and beyond. In Timmons, M., editor, *Oxford Studies in Normative Ethics: Volume 2*. Oxford University Press.
- [Shah and Velleman, 2005] Shah, N. and Velleman, J. D. (2005). Doxastic deliberation. *Philosophical Review*, 114(4):497–534.
- [Sobel, 1976] Sobel, J. H. (1976). Utilitarianism and past and future mistakes. *Nous*, 10(2):195–219.
- [Sobel, 1982] Sobel, J. H. (1982). Utilitarian principles for imperfect agents. *Theoria*, 48(113-26).
- [Thomason, 1981] Thomason, R. H. (1981). Deontic logic and the role of freedom in moral deliberation. In Hilpinen, R., editor, *New Studies in Deontic Logic*, pages 177–86. D. Reidel.
- [Timmerman, 2015] Timmerman, T. (2015). Does scrupulous securitism stand-up to security? two problems for moral securitism and how we might fix them. *Philosophical Studies*, 172:1509–28.
- [van Someren Greve, 2013] van Someren Greve, R. (2013). Objective consequentialism and avoidable imperfections. *Ethical Theory and Moral Practice*, 16(3):481–92.
- [Vessel, 2009] Vessel, J.-P. (2009). Defending a possibilist insight in consequentialist thought. *Philosophical Studies*, 142:183–95.

- [Vessel, 28] Vessel, J.-P. (28). Against securitism, the new breed of actualism in consequentialist thought. *Utilitas*, 2(164-78).
- [von Wright, 1983] von Wright, G. H. (1983). Norms of higher order. *Studia Logica*, 42(2-3):119–27.
- [Vorobej, 2000] Vorobej, M. (2000). Prosaic possibilism. *Philosophical Studies*, 97:131–6.
- [Williams, 1965] Williams, B. (1965). Ethical consistency. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 39:103–24.
- [Woodard, 2008] Woodard, C. (2008). *Reasons, Patterns, and Cooperation*. Routledge.
- [Zimmerman, 1996] Zimmerman, M. (1996). *The Concept of Moral Obligation*. Cambridge University Press.