Some Contemporary Issues about Ought Implies Can: Where Does Kant Fit in?

Zusammenfassung: Die meisten Philosophen stimmen darin überein, dass Kant sich dem Prinzip „Sollte impliziert Können" (OIC) verschrieben hat. Allerdings sind sich nur wenige darüber einig, wie die Bedeutung von OIC zu verstehen ist. Außerhalb der Kant-Wissenschaft gibt es Debatten über die Bedeutung von „sollten", die Bedeutung von „impliziert" und die Bedeutung von „können" in diesem Prinzip. Innerhalb der Kant-Forschung besteht kein Konsens darüber, was Kant zu diesen Themen dachte. In diesem Artikel versuche ich, diese Situation zu verbessern. In Abschnitt I überprüfe ich die Sekundärliteratur zu Kants Engagement für OIC und erkläre, wo es meiner Meinung nach schief geht. In Abschnitt II untersuche ich einige der direkten Textbeweise dafür, dass Kant eine bestimmte Version von OIC zugeschrieben wird. In Abschnitt III lege ich dar, was meiner Meinung nach die wichtigsten doktrinären Gründe dafür sind, diese Version von OIC Kant zuzuschreiben.

Schlüsselwörter: Sollen impliziert Können; Kants Ethik; Kantische Ethik; Korsgaard; Herman; OIC

Abstract: Most philosophers agree that Kant was committed to the principle ′ought implies can′ (OIC). However, few agree on how to understand the meaning of OIC. Outside of Kant scholarship, there are debates about the meaning of ′ought′, the meaning of ′implies′, and the meaning of ′can′ in this principle. Inside Kant scholarship, there is no consensus about where Kant stood on these terms. The present paper tries to go some way toward rectifying this situation. In section I, I review the secondary literature on Kant's commitment to OIC and explain where I think it goes wrong. In section II, I examine some of the direct textual evidence for ascribing a specific version of OIC to Kant. In section III, I set out what I take to be the main doctrinal reasons for ascribing this version of OIC to Kant.

Keywords: ought implies can; Kant′s ethics; Kantian ethics; Korsgaard; Herman; OIC

Most philosophers agree that Kant was committed to the principle ′ought implies can′ (OIC). However, few agree on how to understand the meaning of OIC. Outside of Kant scholarship, there are debates about whether ′ought′ should be understood as *ultima facie* or as *prima facie*; whether ′implies′ should be understood as logical entailment, as metaphysical entailment, as presupposition, or as conversational implicature; and whether ′can′ should be understood as physical ability, psychological ability, general ability, or something else altogether. Inside Kant scholarship, there is no consensus about where Kant stood on these distinctions.

The present paper tries to go some way toward rectifying this situation. In section I, I review the secondary literature on Kant's commitment to OIC and explain where I think it goes wrong. In section II, I examine some of the direct textual evidence for ascribing a specific version of OIC to Kant. In section III, I set out what I take to be the main doctrinal reasons for ascribing this version of OIC to Kant.

I maintain that the direct textual evidence for ascribing OIC to Kant helps to elucidate his position in some of the debates mentioned above. I maintain that the doctrinal reasons for ascribing OIC to Kant also help to triangulate his position. More, they offer an attractive argument for OIC that might be used as a counterweight in current debates about the independent plausibility of OIC. In this way I aim to contribute to both the exegetical and the philosophical disputes about OIC.

1. Why I think current readings of Kant on OIC are mistaken

In this section I review recent work on Kant's commitment to OIC. In particular, I review the work of Christine Korsgaard, Barbara Herman, Markus Kohl, and Robert Stern (in that order). Here are the main mistakes I diagnose:

1. Korsgaard thinks that Kant is concerned exclusively with psychological ability, not with physical ability.
2. Herman and Kohl think that Kant is concerned with general physical ability, not with physical ability in a particular situation.
3. Stern neglects Kant's contrapositive use of OIC.

These are not the only mistakes made by these commentators. But they are the main ones, and they are serious enough, I think, to warrant a reexamination of Kant on OIC. So let me try to make good on these charges.

1. Korsgaard

Korsgaard's account is based on a famous passage from Kant's *Critique of Practical Reason*. In this passage, Kant proposes a thought experiment with two stages. First, he imagines a man who asserts that his lustful inclination for some object is irresistible. Kant says that if a gallows were put up in the vicinity of this object and the man were informed that he would be hanged should he gratify his inclination, then the man would concede his ability to resist the inclination (notwithstanding his original assertion). Second, Kant imagines the same man then being asked whether it would be possible for him to overcome his love of life should a ruler command him, on pain of death, to give false testimony. According to Kant, although the man might not venture to assert that he actually would tell the truth in such a situation, nonetheless, "that it would be possible for him, he must admit without reservation." Kant concludes that the man "judges that he can [do] something on the basis of the fact that he is conscious that he ought [to do] it."

The first stage of this thought experiment does not seem to be morally loaded. Kant is moving from what many would take to be a strong inclination, physical desire, to what many would concede to be an even stronger one, love of life. Thus, the point of the first stage does *not* seem to be that the man in this example judges that he can overcome his lust on the grounds that he *ought* to do so; rather the point seems to be to come up with as strong a desire as possible so that the second stage of the thought experiment is more meaningful: even love of life can be overcome by the moral law, so, clearly, lustful desires can be overcome too. If this is correct, then it is only in the second stage of the thought experiment where OIC is illustrated. The man judges that he is *able* to resist the ruler's command to perjure himself because he judges that he *ought* to do so.

Korsgaard's conclusion from this part of the *Critique of Practical Reason* is that "[w]hen Kant himself advanced [OIC], he meant that it is possible for us to be *motivated* to do what we know we ought to do." She contrasts this with the version of the principle that she says is common in modern discussions: "if you find that it is physically impossible for you to do something, you cannot be obligated to do it." According to Korsgaard, Kant thought that this latter version of OIC is positively mistaken:

> "Kant believed that moral standards, like all rational standards, are essentially human standards, and there is no guarantee that the world will meet them, or make it possible for us to do so."

From this it may be seen that, as Korsgaard reads Kant, ought implies psychological ability, but ought does not imply physical ability. I would like to say three things about this.

First, technically speaking, Korsgaard's Kant is not committed to OIC but rather to what might be called KOIC (i.e., Knowledge-OIC). That is, Korsgaard asserts that if an agent *knows* she ought to D, then she can D. OIC can diverge from KOIC if it is possible for an agent not to know that she ought to D. For example, suppose that I am a physician and I am uncertain whether I ought to kill one of my patients: on the one hand, the patient is terminally ill, in great suffering, and daily requesting to die; on the other hand, I am unsure about whether the patient is in her right mind, and I have conflicting views about physician assisted death. But suppose that, really, I ought to kill this patient. In this situation, OIC would entail that I am able (in some sense) to do so; KOIC would not.

It is unclear to me whether Korsgaard realizes this. One reason for my uncertainty about this lies in the version of OIC she purports to find in the modern debate: if you *find* that D is physically impossible, you cannot have a duty to D. Like the version of OIC Korsgaard ascribes to Kant, the version of OIC Korsgaard ascribes to modern philosophers has an epistemic condition that changes its meaning in a nontrivial way. As will emerge in the next section of this paper, I think *Kant* took there to be epistemic conditions bound up with OIC. But they are not the ones Korsgaard incorporates here, and as far as I know, nobody in the modern debate who subscribes to OIC would say that you might have an obligation to D even though D is physically impossible provided that you have not found out that D is physically impossible. Perhaps charity requires not parsing Korsgaard's text quite so closely.

Second, Korsgaard's explanation of why Kant is concerned only with psychological (rather than physical) ability is based on this conditional:

> If moral standards are essentially human standards, then it might not be physically possible to meet them.

I think that this conditional is mistaken, both philosophically and exegetically. As will emerge in section III of this paper, I think that, for Kant, moral standards are essentially *rational* standards. For Kant, humans are rational beings, but humans are not the only rational beings, nor are humans essentially rational beings, so an essentially rational standard is not an essentially human one. More importantly and as also will emerge in section III of this paper, I think that, for Kant, it must be *physically* possible to meet moral standards *precisely because* moral standards are essentially rational standards (and I see no reason why being essentially human (as opposed to rational) would entail otherwise).

Third and finally, I do not think that the text to which Korsgaard appeals ultimately supports her argument about psychological as opposed to physical ability. I do not want to say that this text positively supports an interpretation of Kant′s OIC as about physical ability. Rather, what I want to say is that this text provides no evidence *against* understanding OIC as about physical ability, and there are other texts that *do* support understanding Kant′s OIC as about physical ability. For example, in the *Critique of Pure Reason* Kant asserts that "the action certainly must be possible under natural conditions if the ought is directed toward it."

2. Herman

Herman's starting point is different from Korsgaard's, and Herman ascribes to Kant a version of OIC that is different from the one Korsgaard ascribes to Kant. Herman starts from ideas about conflicts of duties. She distinguishes two ways that OIC can be understood: (i) as implying a physical ability to carry out some specific action *token* that one is obligated to perform or (ii) as implying a *general* physical ability to carry out an action *type* that one is obligated to perform. For ease of exposition, I shall call these OICi and OICii. According to Herman, OICi has the absurd result that, if one wants to get out of paying a debt, all that one needs to do is squander one's money the day before it is due: in such a scenario the debt no longer can be paid, whence it follows from OICi that the obligation to do so is no longer binding. Thus, Herman advocates for OICii. She then argues that this is how Kant understands OIC and, more, that OICii helps make sense of conflicts of duties because it does not follow from OICii "that a given agent in particular circumstances must be able to satisfy all moral requirements that apply." The idea is that OICii preserves the intuition that conflicts of duty are genuine while also preserving the intuition behind OIC.

However, I think there are three problems with Herman's account. One is that her example about paying a debt misses the mark. Squandering one's money the day before it is due is constitutive of an abrogation of one's obligation to repay a debt. As such, it makes no sense to assert that this is a way of getting out of the token obligation (regardless of the interpretation of OIC). Indeed, this point can be generalized: culpably self-incurred inability (i.e., deliberately making it impossible to fulfill one's obligations) poses no problem for OIC on any interpretation because culpably self-incurred inability is *equivalent* to culpable failure to fulfill an obligation. Thus Herman's *reductio* of OICi does not work and her rejection of OICi is poorly motivated.

The second problem is textual. Herman does not cite a specific passage to ground her reading, so it is difficult to trace it back to its roots. But as I shall try to make clear below (in sections II and III, respectively), there are direct textual and also indirect doctrinal reasons for thinking that, on Kant′s account, ought implies the ability to perform the specific action token that one is obligated to perform (OICi).

The third problem, less relevant for current purposes, is that OICii is not doing the work Herman takes it to be doing in regard to conflicts of duties. To see why, consider a classic example: suppose that I have a duty to go to war to defend my country and a conflicting duty to stay at home to look after my ailing mother. Obviously I cannot do both at once. Herman would say that, because in general I have the ability to to go to war, and because in general I have the ability to stay at home to look after my ailing mother, neither of my duties is nullified by OICii. Thus, according to Herman the mental anguish I feel is warranted, and the theoretical framework she is developing for handling conflicts of duties need not give up on OIC. But exactly the same account is available using OICi. Moreover, note that I do *not* in general have the ability to perform the conjunctive action of going to war *and* staying at home to look after my ailing mother, so Herman's OICii is no better off than OICi as far as conflicts of duties are concerned. The heavy lifting in Herman′s conflict of duties argument is being done by a tacit rejection of the principle of duty agglomeration, and that would work just as well regardless of which version of OIC, OICi or OICii, one subscribes to. I conclude that Herman's ascription to Kant of ′ought implies general physical ability′ (OICii) and her concomitant denial to Kant of ′ought implies specific physical ability′ (OICi) does not withstand critical scrutiny: she does not provide good philosophical or textual reasons for her position, and again, as I shall argue below, there are good philosophical and textual reasons for rejecting it.

3. Kohl

Kohl, by way of contrast with both Korsgaard and Herman, ascribes two versions of OIC to Kant: "I argue that Kant accepts two versions of OIC." One version, like Korsgaard's, concerns psychological ability. The other version, like Herman's, concerns general physical ability.

In my view, Kohl's talk of "two versions of OIC" is misleading at best. Kant is committed only to one version of OIC, and the question is simply what is included in the ability implied by obligation. However, I am going to focus my attention on Kohl's remarks concerning physical ability. My reason for this is that I think Kohl's main error is, like Herman's, to deny that Kant is committed to a version of OIC that includes the physical ability to carry out a specific action token one is obligated to perform:

> Kant's conception of OIC seems to be as follows: a valid prescription that an agent should aim at a certain effect implies that the addressee has the general capacities to produce this effect, but it does not imply that the agent can exercise these capacities in a way that suffices for the actual production of this effect.

Because Kohl identifies his position on physical ability with Herman's, I shall use the same terminology: I shall call the version of OIC Kohl ascribes to Kant concerning physical ability OICii, and I shall call the version of OIC Kohl denies OICi. Kohl has three main arguments, all independent of Herman′s. I am going to assess them in turn.

One of Kohl's arguments is based on rationality. According to Kohl, OICi, unlike OICii, would impugn the rationality of agents on the basis of unforeseeable particularities:

> Since the rationality of our choices cannot be affected by the (for us) unforeseeable vagaries of the empirical world, and since oughts (for Kant) are practical rules that provide conclusive standards for rational choice, these vagaries cannot determine what effects we ought to aim at accomplishing either: oughts must be tailored to the perspective of agents who deliberate from a position of inevitable uncertainty concerning their ultimate success in accomplishing intended effects.

Kohl illustrates this argument with an example. Suppose that Meg has broken legs; in this case, she does not have the general ability to run and, thus, she does not have any duty to save a child who has wandered into the road from oncoming traffic. But if Meg does not have broken legs and has the general ability to run, then "the rationality of her efforts to pull the child off the road is not impugned if these efforts fail because she suffers a cramp in her legs or is hit by a suddenly appearing car."

Kohl's argument here is that if Meg is physically unable to rescue the child because of some unforeseeable condition (like a cramp), then OICi entails that Meg has no corresponding duty of rescue. Kohl maintains that this is absurd because it contradicts the fact that Meg's attempt to rescue, based on her known general ability, is perfectly rational. Thus, OICi, unlike OICii, yields absurd results and should be rejected as inconsistent with Kantian ideas about morality and rationality.

The problem, however, is that Kohl has mistaken the implications of both OICi *and* OICii. Suppose again that Meg is unable to rescue the child because of some unforeseeable condition like a cramp. OICi *does* entail that Meg does not have a duty to perform any action she is physically incapable of performing. But it does *not* entail that she does not have a duty to attempt to do so. That is relevant here because when there is a *prima facie* duty of rescue that Meg incorrectly takes herself to be able to fulfill, OICi leaves intact Meg′s *ultima facie* duty to *attempt* rescue (and thus attempted rescue is rational); it impugns only Meg′s *ultima facie* duty of rescue (and thus failure is not irrational). But OICii leaves intact Meg′s *ultima facie* duty of rescue and thus, if, with Kohl, we affirm OICii and deny OICi, failure *is* irrational although the attempt is not. In other words, if ought entails only general physical capacity and general physical capacity is not impugned by unforeseeable particularities, then an obligation is not nullified by physical incapacity on account of unforeseeable particularities. But from this it follows that (on OICii) Meg's obligation is *not* nullified by her cramp and, thus, her failure to rescue constitutes a failure to fulfill a binding obligation. And from this (conjoined with the Kantian idea, already presupposed by Kohl, that immoral action is irrational) it follows that Meg's failure to rescue on account of her cramp is irrational. It is only with OICi that unforeseeable physical incapacity nullifies an obligation and, thus, it is only with OICi (not OICii) that failure on account of this incapacity ceases to be irrational. To put the point starkly, Kohl's argument shows exactly the opposite of what he wants it to show: unforeseeable and inculpable physical incapacity due to the particularities of a situation coupled with Kantian ideas about morality and rationality gives evidence in favor of, not against, OICi and against, not in favor of, OICii.

Kohl's second argument is based on considerations related to knowledge:

> …on Kant's view every person can know what she morally ought to do (5: 36; 6: 375); empirical circumstances that determine whether our capacities suffice for accomplishing the effects that moral oughts tell us to aim at are typically unknowable for us; thus, the validity of moral oughts cannot depend on those circumstances.

Put schematically, Kohl's argument here is: (I) agents (typically) cannot know whether their capacities suffice for accomplishing the particular (token) effects that moral oughts tell them to aim for; (II) for Kant, every person can know what she morally ought to do; (III) OICi, unlike OICii, entails that agents must know whether their capacities suffice for accomplishing the particular (token) effects that moral oughts tell them to aim for; therefore, (IV) OICi should not be ascribed to Kant.

However, premise (I) in this argument is almost certainly false. I see no reason why agents cannot know that their capacities suffice for accomplishing the particular (token) effects that moral oughts tell them to aim for. For example, if justified true belief is at least sometimes sufficient for knowledge, then an agent who has a justified true belief that she is able successfully to perform some action token on the basis of her general capacities can count as knowing the same just in case nothing untoward is determined to occur, and that is eminently possible.

Kohl's premises (II) and (III) are also objectionable. Regarding premise (II): although in the texts to which Kohl refers in the first sentence of the block quotation above Kant does assert that agents generally are able to determine the deontic status of actions, Kant also seems to leave room for uncertainty. For example, in setting out the casuistical questions in the *Metaphysics of Morals*, Kant does not seem to take the answers to these questions to be immediately obvious, and in his theory of conscience he says explicitly that agents can make mistakes about such judgments. Regarding premise (III): if an agent does not realize that OICi is true, I see no reason why it would entail that an agent must know that her capacities suffice for accomplishing her duties. Further complications arise from the different modalities of these premises and the fact that knowledge is not closed under entailment. From this it may be seen that Kohl's second argument, like his first, fails to withstand critical scrutiny.

Kohl's third argument is based on considerations having to do with luck: "the successful performance of physical action tokens depends on an ineliminable, unforeseeable component of contingency or luck […but] successful compliance with moral prescriptions cannot be a matter of such luck." Kohl's idea here seems to be that (A) if there are obligations to accomplish specific effects, then whether one can carry out these obligations is a matter of luck; (B) OICi entails that there are such obligations whereas OICii does not; (C) Kant thinks that morality is immune to luck; therefore (D) OICi should not be ascribed to Kant.

The flaws in this third argument are similar to the flaws in the second one so I shall be brief. For one thing, whether there are obligations to accomplish specific effects is not germane to the difference between OICi and OICii. So, premise (B) is false. For another, Kant's supposed disavowal of moral luck, although widely maintained, is misguided. So, premise (C) is also false. One way to see this is to look at what might be Kant's most famous thought experiment, the murderer at the door: Kant says that if an agent is untruthful to the murderer and asserts that the victim has gone out, then the agent is partly to blame for the death of the victim should the murderer, as a result of this lie, discover the victim as the victim is hurrying away unbeknownst to the agent. In other words, agents are to be held responsible for the bad effects of their wrongful actions, even if those bad effects are unlucky and result from the intercession of another agent, something Kant reaffirms in the introduction to the *Metaphysics of Morals*. So Kohl's third argument, like his first two, folds under attack.

4. Stern

Stern takes a different approach to Kant and OIC. Like Korsgaard, Stern is concerned with distinguishing Kant's version of OIC from the version used in modern debates. Unlike Korsgaard, Stern does not think this distinction amounts to a difference between psychological and physical ability. Rather, he characterizes Kant as making only a "weak" capacity-expanding inference from ought to can whereas modern philosophers make a "strong" duty-restricting inference from cannot to not-ought:

> [W]hereas the strong conception argues from what we can do to what we ought to do, Kant's weaker conception of 'ought implies can' argues from what we ought to do to what we can do […] It is therefore hardly any surprise that, on close inspection, Kant's position diverges from the current one.

But there are at least two problems with this.

The first problem has to do with Stern′s distinction between strong and weak versions of OIC. As I argue in the next section of this paper, close inspection of Kant′s texts reveals that he subscribed to ought *logically entails* can, and if this is correct, then the strong and weak versions of OIC are inseparable. Moreover, other versions of OIC, like ′ought conversationally implicates can′ and ′ought presupposes can′, would also create problems for Stern. Conversational implicature does block the duty-restricting use of OIC (because implicatures can be cancelled), but it also blocks the capacity-expanding use of OIC. If, by way of contrast, ought presupposes can, then inability to D does not show that it is *false* to assert an obligation to D; it shows that it is neither true nor false to assert an obligation to D, whence it follows that it is *true* to assert the absence of an obligation to D. So if ought presupposes can, we still get a version of the duty-restricting ′cannot implies not-ought′.

Elsewhere in his article Stern explains the strong/weak distinction in terms of a difference between the moral law (*simpliciter*) and the moral law as it applies to us. Kant does say that the Supreme Law of Morality manifests as an ought only for imperfectly rational beings like us, not for perfectly rational beings. But, problematically for Stern, Kant maintains that this is so because imperfectly rational beings must be *constrained* to follow the moral law:

> It [the moral law] is thus not limited merely to humans, but rather extends to all finite being which has reason and will, indeed includes within [its bounds] even the infinite being as supreme intelligence. In the first case however the law has the form of an imperative, because one can presuppose in the former indeed as rational being a pure will, but as being affected by needs and sensible motives no holy will, i.e. one such that it would be capable of no maxims that are contrary to the moral law.

In other words, the difference between the moral law and imperatives, on Kant′s account, is that the latter are directed at beings with the ability to do otherwise. So this attempt to uphold the strong/weak distinction fares no better.

However, the second problem with Stern′s interpretation is deeper. I concede that it is the capacity-expanding use of OIC that is most widely associated with Kant: the passage used by Korsgaard (examined above), for example, is quite famous, as is Kant's highest good argument (examined below), and part of what these are famous for is their incorporation of OIC in its capacity-expanding form. But Kant makes use of OIC in its duty-restricting form all the same. For example, consider the following three passages, the first from the *Metaphysics of Morals*, the second from the *Religion within the Boundaries of Mere Reason*, the third from *On the Common Saying*:

> Love is a thing of sensibility, not of the will, and I cannot love because I will, much less however because I should […] hence a duty to love is an absurdity.

> [T]he elevation of such a holiness over every frailty of human nature would be […] in the way of [i.e., an obstacle to] the practical application of the idea of this being for our imitation […] this distance from natural humans thereby would become again so infinitely great that the former divine human could not be set up for the latter any more as an *example*.

> [I]t would not be [a] duty to aim at a certain effect of our will if this were not also possible in experience (which might now be thought as completed, or as always approaching completion).

In the first of these passages, Kant asserts that there can be no duty to love because agents are unable to love at will. In the second, Kant asserts that a morally perfect being cannot be a moral example because such heights are, for natural humans, unattainable. In the third, Kant asserts that there can be no duty to aim at some effect in experience if that effect is not also possible. In all of these passages, Kant is using OIC in its contrapositive form, and these are not the only ones there are: in the *Critique of Practical Reason* Kant goes so far as to suggest that, if there is a duty that cannot be carried out, then the moral law would cease to be ″valid.″ Stern is aware of and tries to defuse this passage by noting that ″Kant does not need″ to make such a strong claim. But whether Kant needs to make such a claim is irrelevant in the face of the fact that he does so.

5. Interim Conclusion

Having explained why I think that Korsgaard, Herman, Kohl, and Stern go wrong in their own separate ways, I want to conclude this section by diagnosing a common mistake: all of these commentators would deny that, for Kant, being physically incapable to D in a specific instance entails that there is no duty to D in that instance. Indeed, not only do I think that this denial is mistaken, but, more, I think Kant takes ought to entail much more than physical ability. I think that there is both direct textual and indirect doctrinal support for this, and, as noted above, I think that the latter can provide welcome support to those engaged in the modern debate about the plausibility of OIC.

II. Direct textual support

I take it that the question now is not whether Kant was committed to OIC but rather how to understand the version of OIC to which he was committed. Accordingly I am going to focus on seven key passages that provide direct textual support for my thesis that Kant is committed to a very specific version of OIC: *ultima facie* ought logically entails (i) real possibility of performance of specific action token, (ii) epistemic warrant for believing in conditions necessary for real possibility, and (iii) epistemic ability (i.e., know-how). Before turning to these seven passages in order to defend and explicate this thesis, however, I want to say something about Kant's highest good argument.

As noted above, Kant's highest good argument involves one of his most famous uses of OIC. The highest good is a world in which all agents are supremely virtuous and happiness is distributed according to virtue. According to Kant, the highest good, "insofar as it is attainable, is also duty, and conversely, if it is duty, also must be attainable." Kant famously uses this to derive an epistemic warrant for belief in God and immortality, two conditions for the real possibility of the highest good. His reasoning is something like this: (i) there is a duty to promote the highest good; therefore (ii) it must be really possible (in a sense to be explicated momentarily) to promote the highest good; but (iii) the highest good is really possible only if we are immortal and God exists; therefore (iv) we have an epistemic warrant to believe we are immortal and God exists.

From this, I would like to draw two important lessons. First, Kant is not committed merely to ′ought implies logical possibility′ or even the stronger ′ought implies physical possibility′: Kant is committed to ′ought implies real possibility′, or ′ought implies really can′. That is, on Kant's account, if I ought to D, then D must be logically and physically possible *and it must have some ground in reality*. The highest good is logically and physically possible on its own. But, according to Kant, it is a real possibility only if it can be grounded in some real things (immortal souls and God). Thus, on Kant′s account, God and immortality may be presumed real precisely because they ground the possibility of the highest good and therefore the bindingness of the duty to promote the highest good.

Second, Kant's version of OIC is bound up with epistemic conditions. Some care is needed here because Kant does *not* infer, in his highest good argument, that agents *ought* to believe in God and immortality. So Kant might not be committed, for example, to ′ought implies ought to believe really can′. Rather, Kant infers that agents *may* believe in God and immortality. So Kant *does* seem to be committed, on the basis of the highest good argument, to something like, ′ought implies *permission* to believe really can′. I shall return to this momentarily. For now, I would like to interrogate the following 7 passages to help figure out how Kant understood OIC:

1. The subjective grounds and calculations of imputation are: the intention. The knowledge. The capacity. The readiness. The opportunity.
2. Pure reason thus contains, indeed not in its speculative but nevertheless in a certain practical, namely the moral, use, principles of the possibility of experience, namely such actions which, according to moral prescripts, could be met with in the history of humans. For since it is prescribed that such [actions] ought to happen, so they must also be able to happen […]
3. [T]he concept of duty […] contains that of a good will, although under certain subjective limitations and hindrances […]
4. Duty is the necessity of an action from respect for the law.
5. [T]his ought is actually a will which is valid for every rational being under the condition that reason in him would be by it practical without hindrance […]
6. [D]uty charges an agent unconditionally: he *ought* to remain true to himself; and from this he rightfully *infers*: he must also *be able* to do so, and his will is thus free […] freedom, according to which the action as well as its opposite must be, in the instant of occurrence, within the power of the subject.
7. When I *ought* to do something, *then* I must also *be able* [to do] it and what is unpardonably behooved to me must also be possible to me to attain.

Although I shall not be making any arguments here about the evolution of Kant's views, I have ordered these passages chronologically. The first is from Kant's handwritten *Nachlass*, an unpublished reflection dated to 1776-1778; the second is from the *Critique of Pure Reason*; the third, fourth, and fifth are from the *Groundwork to the Metaphysics of Morals*; the sixth is from the *Religion within the Boundaries of mere Reason*; and the seventh is from the first convolute of the *Opus Postumum*. While I realize that these 7 passages are unconventional ones to which to appeal in a discussion of Kant on OIC, I nonetheless do so because I think they are particularly suited to my purposes: dialing down to the specifics of Kant's understanding of this principle against the backdrop of a common understanding, from the foregoing, that he was committed to it.

Passages 3, 4 , and 5 are revelatory for the interpretation of ′implies′ in Kant's OIC. In 3, Kant remarks that the concept of duty *contains* that of a good will. This involves an appeal to Kant's understanding of logic and, in particular, concept containment. The idea is that each concept has various marks associated with it. For example, the concept ′vertebrate′ contains the mark ′having a backbone′ and, thus, on Kant′s account, having a backbone is logically entailed by being vertebrate. So in 3, Kant asserts that the set of marks analytic to the concept of duty includes all the marks associated with a good will. That is precisely why, in 5, Kant asserts that, for a being that is perfectly rational, one in whom reason is ″practical without hindrance″ or, in the words of 3, one whose good will is *not* under ″subjective limitation and hindrances,″ "ought is actually a will.″ Similarly, in 4, Kant *defines* duty as ″the moral necessity of an action from respect for the [moral] law.″ The necessity here is on account of the subjective limitations and hindrances that we, as imperfect rational beings, labor under. So, according to Kant, if the moral law says that I should D and if I have a perfectly good will, I shall D; but because I do not have a perfectly good will, I might not D, and for that reason I am subject to constraint. From this it may be seen that Kant takes the ′implies′ in OIC to be strict logical entailment: the concept of ′duty′ or ′ought′ *contains* the doing of the action in question, and it is only our subjective limitations and hindrances that trip us up.

Recall the debates gestured to in the introduction of this paper. Some modern philosophers who defend OIC think that it should not be interpreted in terms of logical entailment. For example, consider Frances Howard-Snyder:

> I shall interpret the third term in OIC, "implies," roughly as stating a metaphysically necessary connection […] I do not mean to suggest that the connection is one that every competent speaker of the language will agree is obvious. OIC cannot be refuted by the mere fact that some have questioned it.

However, this move away from logical entailment and towards metaphysically necessary connection seems precipitate to me. Many logical entailments are far from obvious, and competent speakers need not have grasped every nuance of every concept in the language. So if Kant was wrong to subscribe to ′ought logically entails can′, I do not think that this argument shows why.

Passages 3, 4, 6, and 7 bear on Kant's understanding of ′ought′ in OIC. In particular, I think that these passages indicate that Kant was more interested in the all-things-considered ′*ultima facie* ought implies can′ (UFOIC) than the some-things-considered ′*prima facie* ought implies can′ (PFOIC). Passage 4 defines duty in terms of what the law *necessitates* me to do, not what the law gives me reason, but perhaps not overriding reason, to do. So 4 is about *ultima facie* duty. Moreover, the proximity of 3 to 4 in the text suggests that the containment claim in 3 is about *ultima facie* duty as well: it is *ultima facie* duty that contains the concept of a good will. Passage 6 is even more telling insofar as it says that the action that an agent ought to perform must be possible for her ″in the instant of occurrence.″ Passage 7, in contrasting what I ought to do with what I am ″unpardonably behooved″ to do and inferring ability from both, might suggest that Kant is committed to both PFOIC *and* UFOIC. That is, some might think that the first ′ought′ in 7 is *prima facie* ought whereas what I am ′unpardonably behooved′ to do is what I *ultima facie* ought to do. But either way, 7 is, I think, good evidence of Kant′s commitment to UFOIC.

Passages 1, 2, 5, 6, and 7 bear on the question explored above in discussion of Herman and Kohl. I think that these passages give evidence in favor of interpreting the physical ability (and real possibility) implied by obligation in terms of a specific action token rather than in terms of general capacity. Passage 1 even uses modern terminology that is deployed when philosophers want to distinguish between the ability to perform a specific action and a more general ability: in 1 Kant says that both ″the capacity″ to carry out an obligation as well as ″the opportunity″ to do so are implied by ″ought″, whence it follows that this capacity cannot be masked (e.g., by a cramp, as in Kohl′s example). Passage 5, in stating that an imperfectly rational being would do that which he ought if reason had full, unhindered control over his will, would make little sense if the action that an agent was obliged to perform was physically impossible in that instance. Similarly for 2 and 7: passage 2 says that the principles of morality are ″principles of the possibility of experience,″ very telling language for Kant in particular, and 7 says that agents must be able ″to *attain*″ that which is ″unpardonably behooved″ to them, their *ultima facie* obligations. And as noted in the previous paragraph, passage 6 says that an obligation must be ″within the power of the subject″ in the very instant in which it is to be performed. This is, I think, a clear disconfirmation of the claim made by Herman and Kohl that, on Kant′s account, ought implies only general ability.

I want to make two last points about these quotations. First, Passage 6 bears on the question explored in my discussion of Stern. Recall that Stern attempts to distinguish the Supreme Law of Morality from the Categorical Imperative: the former applies to all rational beings whereas the latter applies only to imperfectly rational beings like humans, and Stern seems to suggest that this means that the former does not imply can whereas the latter does. But as I pointed out above, this is mistaken. On Kant′s account, what this means is that the latter implies ′able not to′ whereas the former does not. I take passage 6 to confirm this interpretation, for in passage 6 Kant says that ought implies ″freedom, according to which the action *as well as its opposite* must be, in the instant of occurrence, within the power of the subject." In other words, according to Kant, ought implies both can and able not to.

The last thing I would like to say in this section concerns know-how. This hearkens back to the remarks made above in the discussion of Korsgaard about epistemic conditions. Passages 1 and 5 both support ascribing to Kant a robust ability that includes not merely real possibility of a specific action token but also the knowledge of how to perform it. Passage 1 says that ″the knowledge″ of how to carry out one′s duty is implied by ought, and 5 would make no sense if the agent did not know how to fulfill her obligation: the counterfactual world Kant is envisioning in passage 5, a world in which our reason is ″practical without hindrance″ and so the ″ought″ of duty ″is actually a will″, is a world in which *motivational* hindrances to duty are not present; it is *not* a world in which epistemic (or physical) hindrances to duty are not present.

To illustrate, it might be physically possible for an agent to pay a debt, and she might have the financial and motivational wherewithal to do so. But if the only means in a given context is to use a complicated electronic payment system that she is, for whatever reason, unable to figure out, and if ought implies not merely physical ability but (what might be called) epistemic ability, then her duty to pay back the debt is no longer binding. Now, if the agent merely does not know how to use the payment system (as a matter of actual fact), then her duty to pay the debt might generate a further duty to learn how to use the payment system. In that kind of scenario the bindingness of the debt remains intact. But that is not the kind of scenario I am describing. As I am imagining the scenario, the agent is unable to attain this knowledge, perhaps because there are time constraints on when the debt must be paid, and it is important to bear that in mind when thinking about the plausibility of Kant's principle. We might imagine the agent having agreed to give the person who is buying her house a credit at closing. To do so, she brings her checkbook, only to be told that the buyer requires the transfer to be made via Venmo, which she never has used before. In this sort of scenario, the agent′s obligation to credit the buyer at the closing is, at least on Kant′s version of OIC, nullified: the buyer can accept the check or give the seller sufficient time to get comfortable with Venmo. The seller′s obligation to pay the credit does not go up in smoke. Rather, the moral landscape has shifted on account of her inability to master a new payment method on short notice.

On the basis of this, I conclude that, for Kant, *ultima facie* ought logically entails real possibility of fulfillment of specific action token, epistemic warrant for believing conditions necessary for real possibility, and epistemic ability, quite a strong principle indeed. In the next and final section, I shall explain what I take to be Kant's reasons for subscribing to this.

Section 3. Indirect doctrinal support

In this section, I am going to discuss the indirect doctrinal support for ascribing to Kant the specific version of OIC discussed above. As I understand Kant, his argument for subscribing to OIC may be reconstructed as follows:

1. If an agent morally ought to D, then she has most reason to D.
2. If an agent has most reason to D, then in the world in which she is governed by reason, she will D.
3. If there is a possible world in which an agent will D, then it is possible for her to D.
4. Therefore, if an agent ought to D, then it is possible for her to D, i.e., OIC is true.

Now the language in this argument is entirely different from the language in the passages in the previous section. So I want to say something about each of the premises in this argument in order to tie them back to Kant. I also want to try to make good on a promissory note from the introduction: I want to explain why I think Kant′s argument should play a role in current debates about the plausibility of OIC.

Premise 1 encapsulates the well-known Kantian bridge principle, bridging morality and rationality, that says that morality is a form of rationality. This principle is found in passage 5 (see above, at note 42) and, arguably, passage 4. In passage 5, Kant says that rational beings, if they are acting rationally, always will conform to the moral law. That is precisely why ″ought is actually a will […] under the condition that reason in him would be by it practical without hindrance.″ Similarly, the definition of duty in passage 4 as ″the necessity of an action from respect for the law″ reflects the idea that moral oughts are categorical, overriding all other oughts (including councils of prudence). One reason this is notable is that, although some philosophers reject Kant′s bridge principle, my way of formulating the Kantian argument for OIC makes evident that these philosophers still would be able to subscribe to a version of OIC; it is just that their version of OIC would not be *moral* ought implies can. Rather, it would be *rational* ought implies can.

Premise 2, like premise 1, is built from the passage 5 claim about an ought being a will when an agent is governed by reason. But premise 2 also reflects passage 3, which says that the concept of duty contains that of a good will ″under certain subjective limitations and hindrances.″ A perfectly good will, one not subject to these limitations and hindrances, will do that which the moral law commands straightaway: on Kant′s account, it is because of these limitations and hindrances that the moral law manifests as duty to us. Duty is a form of constraint, required for (but only for) imperfectly rational beings. According to Kant, the limitations and hindrances that prevent an agent from doing what she ought are *motivational*. Evil, according to Kant, is not to be understood in terms of ignorance or false beliefs, nor is it to be understood in terms of misperception, misjudgment, or poorly functioning faculties more generally. Similarly, Kant thinks that evil is not a mere failure to promote the right results or a failure to comply with social norms. Evil, on Kant′s account, is willful: it involves adopting a maxim, acting *in accordance with the representation of a principle* that cannot be willed at the same time as a universal law or that involves failing to treat humanity at the same time as an end. This, of course, is bound up with an internalist account of morality and (apropos of premise 1) rationality. But the point for present purposes is that it is because of how Kant explains the difference between an imperfectly and a perfectly rational being (namely, in terms of subjective hindrances and limitations) that his version of OIC implies so much. Let me unpack this.

Because evil is not merely a matter of ignorance or false beliefs but rather a matter of maxims, if an agent′s action fails to conform to the moral law, this cannot be on account of her ignorance about what the moral law requires of her or how to carry it out. This does not mean that agents have infallible and immediate knowledge of the requirements of morality. On the contrary, it means that, when agents act in accordance with their best but fallible judgments about the deontic status of a given action, they will have done all that morality can require of them. Kant embraces this explicitly in the *Metaphysics of Morals* when, shortly after noting that ″in the objective judgment whether something is a duty or not one can well err,″ he asserts that ″[i]f, however, someone is conscious of having acted according to conscience, then, as far as guilt [*Schuld*] or innocence [*Unschuld*] is concerned, nothing more can be demanded of him.″ It is from this, I think, that Kant can make an argument for the claim that ought implies epistemic ability: if an agent fails to do that which she morally owes [*schulden*] merely on account of ignorance, she is in the clear. This also goes to show one of the strengths of this argument. Premise 2 can be interpreted in different ways. If, with Kant, we maintain that what hinders imperfectly rational agents from acting perfectly rationally is motivational, then, as I have been arguing, ought will imply epistemic conditions of fulfillment. If, however, we adopt a more expansive account of the difference between imperfectly and perfectly rational beings, then we will wind up with a different version of ought implies can.

The third premise makes evident that I am portraying this argument in terms of a possible worlds account of modality, notable because there is debate about whether that is the appropriate interpretation of modality according to Kant. However, I do not think that this is an essential part of the argument: I think the argument easily could be reworked in accordance with an alternative account of modality to get to the same conclusion. For example, premise 2 could be retooled as: if an agent has most reason to D, then in my conception of her as governed by reason, she will D, and premise 3 would be, accordingly: if I can conceive of an agent as performing D, then it is possible for her to D.

Reframing the third premise in this way makes clear, I think, that the lynchpin of this argument is the claim that it is always possible for an agent to be governed by reason or, in Kant′s terms, it is always possible for reason to be practical. Obviously I cannot give a full defense of this claim here (such a defense would would have to delve into Kant′s ideas about freedom). But one of the merits of this discussion is that it shows, I think, that Kant′s version of OIC is part of a much larger vision of morality, evil, and rationality. Although I have tried to distill the core of his argument into three premises that are separable from this larger vision, it seems to me that trying to assess OIC in a vacuum, divorced from these larger conceptual questions about the nature of evil and what it means to be imperfectly rational, as is done especially in modern exphi, is, at best, going to give a distorted picture of the truth. Kant's contribution to the modern debate about OIC thus should be a welcome supplement, one that ought to be engaged with by all participants—and one that therefore can be so engaged.

Conclusion

In this paper, I have tried to go some way toward resolving the complex knot of issues surrounding OIC both within and outside of Kant scholarship. I began by canvassing recent work on Kant′s commitment to OIC. I diagnosed various mistakes in this literature, culminating in one common problem: none of them can accommodate the Kantian implication of OIC that inability in a specific instance nullifies any corresponding obligation. I then reviewed some of the direct textual evidence for ascribing OIC to Kant. I argued that Kant is committed to a very specific version of OIC: *ultima facie* ought logically entails real possibility of fulfillment of specific action token, epistemic warrant for believing conditions necessary for real possibility, and epistemic ability. In the final section, I reconstructed a Kantian argument for OIC, explaining where I find evidence for this argument in Kant as well as why I think it would be a welcome addition to current debates.