An Empirical Refutation of 'Ought' Implies 'Can'

PAUL HENNE, VLADIMIR CHITUC, FELIPE DE BRIGARD AND
WALTER SINNOTT-ARMSTRONG

The famous principle that 'ought' implies 'can' claims:

   (OIC) For any agent, J, and act, K, if J ought to do K, then J can do K.

   Many philosophers, including Kant and Smith (1933: 473), Sidgwick (1884: 33), Moore (1922: 317) and Parfit (1984: 15) accept this universal generalization. For most of them, OIC is true not only universally but also analytically (Zimmerman 1996: 79) or conceptually (Vranas 2007: 171). Indeed, it is difficult to see why OIC would be true if not by virtue of the concepts expressed by 'ought' and 'can'.

   Despite its popularity, a few contrarians, including Stocker (1971), White (1975, 1979), Sinnott-Armstrong (1984, 1985) and Graham (2011), reject OIC. Although these skeptics typically admit some related principle – such as 'ought' conversationally implies 'can' (Sinnott-Armstrong 1984) – they often deny that OIC holds universally on the basis of their own judgements about thought experiments presented as counter-examples. Nonetheless, whether their judgements about such thought experiments reflect people's actual concepts and linguistic practices remains unclear.

   We think that new empirical evidence elucidates this issue. After discussing our methodology (§1), we will describe empirical evidence that suggests that people often do accept the counter-examples and reject OIC (§2). Then we will consider responses (§3) before concluding that OIC is not analytically or conceptually true (§4).

1. Methodology

One test for analyticity or conceptual entailment is simple. If one claim analytically or conceptually entails another, then competent speakers in good epistemic positions should consistently deny the first when they know that the second is false. Consider the claim that 'bachelor' analytically entails 'male'. With this entailment, we can predict that competent speakers who know that Pat is not male should deny 'Pat is a bachelor'. If, however, most competent speakers were to accept, 'Pat is not male, but she is a bachelor,' there would be strong evidence that 'bachelor' does not analytically or conceptually entail 'male' (cf. Stich and Weinberg 2001). This schema is a test for analytic or conceptual entailment.

   Likewise, if 'ought' analytically or conceptually entails 'can' – as many philosophers claim – then competent speakers should deny that an agent

ought to do an act when they understand that the agent cannot do the act. If competent speakers deny that an agent ought to do an act when they understand that the agent cannot do the act, then this finding is consistent with OIC. If, however, competent speakers assert that an agent ought to do an act when they understand that the agent cannot do the act, then this finding is inconsistent with the claim that OIC is analytically or conceptually entailed. So, do they?

2. Some experiments

Recent experimental work suggests that ordinary people do not make judgements in accordance with OIC. In one recent study, Chituc et al. (An empirical refutation of ought implies can. (under review)) asked participants to judge whether an agent ought to keep a promise (to meet a friend for lunch at noon) that he is unable to keep either because he chose not to leave in time (high-blame condition) or because his car breaks down (low-blame condition). Participants read the following vignettes (the first paragraph is constant between the two conditions):

> Adams promises to meet his friend Brown for lunch at noon today. It takes Adams thirty minutes to drive from his house to the place where they plan to eat lunch together.
>
> *Low blame:* Adams leaves his house at eleven thirty. However, fifteen minutes after leaving, Adams car breaks down unexpectedly. Because his car is not working at that time, Adams cannot meet his friend Brown at noon, as he promised.
>
> *High blame:* Adams decides that he does not want to have lunch with Brown after all, so he stays at his house until eleven forty-five. Because of where he is at that time, Adams cannot meet his friend Brown at noon, as he promised.

After reading both vignettes in random order, participants were asked, 'Do you agree or disagree with the following statement: At 11.45, it is still true that Adams ought to meet Brown at noon.' Participants answered on a scale from −50 ('completely disagree') to 50 ('completely agree'), with 0 being 'neither agree nor disagree'. Participants' responses significantly differed between the two conditions. In the low-blame condition, most participants (68%) – as defenders of OIC might expect – denied that the agent ought to keep his promise – that is, their answers were significantly below the midpoint. In the high-blame condition, however, the majority of participants (60%) judged that the agent ought to keep his promise even when they knew he was unable to do it; in other words, most participants asserted 'ought' without 'can'.

In a follow-up experiment employing similar low-blame vignettes – where judgements are less likely to be distorted by blame (cf. Alicke 2000) – Chituc et al. (2016) collected participants' judgements of 'ought', 'can' and 'blame'. Against the prediction of OIC defenders, they found no significant correlation between participants' 'ought' judgements and their 'can' judgements but they did find a significant correlation between their 'ought' judgements and their 'blame' judgements. In a third experiment, Chituc et al. presented participants with similar vignettes in which agents are unable to perform an act but this time they manipulated both whether or not the agent was blameworthy for failing to perform the act and whether or not the agent had a moral obligation to perform the act (either by making a promise or by expressing a desire to do something). Consistent with the previous experiments, participants judged that the agent ought to perform the act even if he cannot but they made this judgement only when the agent was blameworthy for not being able to perform an act they were morally obligated to perform. The researchers found no such interaction in conditions where the agent was not blameworthy or where there was no moral obligation. As such, the results of this third experiment strongly suggest that participants understand the concepts of 'ought' and 'can', and they can apply them discriminatively across moral and non-moral contexts and across fault and non-fault situations.

Buckwalter and Turri (2015) reported consistent results from a series of studies examining judgements of obligation and ability in an array of contexts. In their forced-choice experiments, they varied the type, duration and scope of the inability; the vocabulary describing the moral obligation; the consequences of inaction; and the type of obligation (for instance, *moral* or *legal*). Yet participants repeatedly attributed obligations to agents who were unable to perform the relevant acts. In Experiment 1, for instance, participants read short vignettes about an agent, Walter, who promised to pick up a friend, Brown, at the airport but was physically unable to do so (Buckwalter and Turri 2015). Participants were then required to choose one of four statements about Walter's obligation and ability to pick up Brown – two denying his obligation (one denying his ability and the other asserting it) and two asserting his obligation (one denying his ability and the other asserting it). Eighty percent of participants who read this vignette selected the following statement: 'Walter is obligated to pick up Brown at the airport but Walter is not physically able to do so'. Buckwalter and Turri confirm these results throughout their article, and they also find that participants' judgements differ in non-moral contexts. This work (cf. Mizrahi 2015) accords with the findings of Chituc et al. (2016)

The results of these two studies challenge the claim that OIC is analytically or conceptually true. If 'ought' judgements entail 'can' judgements, we would expect those judgements to correlate – *ceteris paribus*. Chituc et al. (2016), however, found that most participants readily say that an agent ought to do what they know the agent cannot do. Notably, when participants were

prompted to explain their responses, one participant commented, 'Adams won't be ABLE to but he still OUGHT to' (capitals in participant response). Assuming that these speakers are representative and competent, these findings suggest that most competent speakers do not deny that an agent ought to do an act whenever they understand that the agent cannot do the act. Thus, this finding is inconsistent with the claim that OIC is analytically or conceptually true.

These findings may also explain of why most philosophers accept OIC: they seem to conflate 'ought' with 'blame'. Specifically, if philosophers assume (1) agents are not blameworthy for failing to do what they cannot do and (2) agents ought to do only what they would be blameworthy for failing to do, then these philosophers can conclude that 'ought' implies 'can'. This argument is explicit when Brown writes, 'Does ought imply can? Surely it does. For we do not hold a person to blame for not doing something he was unable to do' (Brown 1977: 206; cf. Parfit 1984: 36 and Copp 1997: 445). The studies by Chituc et al. (2016), however, suggest that people reject Assumption (1) in cases where an agent is blameworthy for being unable to do what they claim he ought to do. Overlooking these cases where ascriptions of blame modulate 'ought' judgements may explain the broad acceptance of OIC.

3. Objections

In this section, we review and reject several possible replies by defenders of OIC.

3.1 Competence or performance?

One reply to the empirical findings is that they reveal performance errors rather than linguistic competence (cf. Nichols and Knobe 2007: 671). Perhaps participants believed that the agent could keep his promise because they misinterpreted or forgot the presented scenario. There is, however, no evidence of such an error occurring. Chituc et al. (2016) asked participants whether the agent was able to do the act either right before or right after they asked whether the agent ought to do the act, and they asked participants whether the agent ought to do the act at precisely the same specified time when participants agreed the agent could not do the act. Moreover, they (Chituc et al., Experiment 2) separately analysed the results excluding those speakers who may have been making mistakes – they dropped participants who said that the agent could perform an act that he could not perform – but the results did not change. Hence, it is unlikely that participants were confused about the agent's ability.

A similar reply is that participants were in poor epistemic conditions (cf. Williamson 2010), so improving epistemic conditions – for instance, by allowing participants to see contrasting vignettes – may reduce some asymmetric judgements (cf. Pinillos et al. 2011). Chituc et al.'s results, however,

held even when participants saw both vignettes in a within-subjects design, and there were no ordering effects (Chituc et al., Experiment 1). Moreover, there is evidence that participants discriminated very finely between concepts, so they were in quite good epistemic conditions; both in Chituc et al. (Experiment 3) and in Buckwalter and Turri (Experiment 7), for instance, participants responded differently to moral and non-moral contexts. Nevertheless, while it is impossible to eliminate all performance errors, there is no reason to suspect such an error in the results of either study.

3.2 Which concept?

Defenders of OIC might argue instead that participants in the experiment use a different concept of 'ought' than the philosophically refined one that implies 'can'. On this defence, the participants express $ought_F$, which does not entail can and is distinct from $ought_P$ that moral philosophers express.

There are two difficulties with this defence. First, ordinary speakers like the participants in these studies presumably use ordinary concepts. If philosophers use a distinct concept when they endorse OIC, then they need to disclose it, describe the refined concept, and explain why it is better than the ordinary one. Philosophers, however, rarely do so; usually, they describe *ought* as the concept that we ordinarily use (e.g. Moore 1922: 317). Of course, it is still open to philosophers to defend their alternative $ought_P$ concept – although such a defence is absent from the literature.

Even if philosophers did defend a special concept that is distinct from the ordinary concept, a second problem would arise. Specifically, it is unclear how the theories that use some concept like $ought_P$ address the moral issues that interest non-philosophers who use the ordinary concepts like $ought_F$. That conceptual disconnect prevents philosophical theories from informing ordinary moral issues, which seems to be the aim of many of these moral theories. Philosophers, of course, can continue to hold this distinct concept while rejecting the empirical data. But the defence here seems ad hoc – only for the sake of preserving some form of OIC – and robs their position of its interest to most people.

3.3 Modus tollens

Defenders of OIC might also appeal to a modus tollens argument. Imagine that Brit promises to meet Dan for a beer. It seems that she ought to keep her promise. But suppose that Brit is in an automobile accident on the way to meet Dan. Now, Brit cannot meet Dan. In such a case, most people retract the judgement that Brit ought to meet Dan. Cases like this one suggest the generalization:

(OIC*) For any agent, J, and act, K, if agent J cannot do act K, it is not the case that J ought to K.

(OIC*) entails (OIC).

The problem with this reply is that speakers who withdraw their 'ought' judgements in such cases may assume more than just the agent's inability to perform that act. They may assume not only that Brit is unable to keep her promise but also that her inability is not her fault. Even if the accident was due to Brit's negligence, they may assume that Brit did not intentionally cause the accident to evade her meeting with Dan. Subsequently, they could retract their 'ought' judgement not only simply because Brit cannot meet Dan but also because Brit is not blameworthy for her inability to meet him. If so, exculpation rather than inability explains their retraction of the 'ought' judgement. And exculpation, as we discussed, does not always follow from inability. Hence, that retraction fails to support OIC even if it suggests that 'ought' implies 'blame', which accords with Chituc et al.'s results.

4. Conclusion

The results of Chituc et al. (2016) and Buckwalter and Turri (2015) undermine OIC; given these results, OIC is not true analytically or conceptually. And if OIC not true analytically or conceptually, it is difficult to see why it is true at all. Furthermore, these results are not the consequence of performance errors or of epistemic limitations, and they cannot be avoided by stipulating a kind of 'ought' that does imply 'can'.

Philosophers still might try to restrict OIC to special cases. For example, the vignettes in Chituc et al. used third-person 'ought' judgements about the past, so 'ought' still might imply 'can' in first-person 'ought' judgements about the present or about the future (cf. Streumer 2003: 222–24). Although this tensed version of OIC may be true, the restricted generalization cannot salvage the traditional version of OIC that is supposed to be true by virtue of the meanings of the terms 'ought' and 'can'. This salvaging of OIC fails because the meanings of those terms do not change when they are used in the first-person as opposed to the third-person or when they are used about the present as opposed to the past. To restrict OIC to special cases is, in effect, to abandon the traditional universal principle.

Hence, we propose the end of OIC as it is traditionally understood. At least, our argument shifts the burden proof onto those who currently accept it. Thus, philosophers must either find a better argument to support OIC or abandon it.[1] Until they meet this burden, philosophers should stop assuming OIC unconditionally. And without this assumption, they will need to rethink

---

1   Many readers will note here that there are other arguments in favour of OIC that do not rely on analyticity or on conceptual entailment. For example, there are arguments that rely on views of deliberation (Hare 1963: 51–61; Streumer 2007: 365). These arguments, as well as some others, may successfully avoid being affected by the empirical evidence discussed here. Nonetheless, these arguments fail in other ways, which the authors aim to show in a companion article.

a vast array of positions – including many theories of moral responsibility – that appeal to OIC (cf. Driver 1983).[2]

Duke University
Paul.Henne@Duke.edu,
Vladimir.Chituc@Duke.edu,
Felipe.DeBrigard@Duke.edu,
Walter.Sinnott-Armstrong@Duke.edu

References

Alicke, M. 2000. Culpable control and the psychology of blame. *Psychological Bulletin* 126: 556–74.

Brown, J. 1977. Moral theory and the ought-can principle. *Mind* 86: 206–23.

Buckwalter, W. and J. Turri. 2015. Obligation and inability in moral judgment. *PLoS One* 10: e0136589.

Chituc, V., P. Henne, W. Sinnott-Armstrong and F. De Brigard. 2016. Blame, Not Ability, Impacts Moral "Ought" Judgments for Impossible Actions: Toward an empirical refutation of "Ought" implies "Can". *Cognition* 150: 20–25.

Copp, D. 1997. Defending the principle of alternative possibilities: blameworthiness and moral responsibility. *Nous* 31: 441–56.

Driver, J. 1983. Promises, obligations, and abilities. *Philosophical Studies* 44: 221–23.

Graham, P. 2011. 'Ought' and ability. *The Philosophical Review* 120: 337–82.

Hare, R.M. 1963. *Freedom and Reason*. New York: Clarendon Press.

Kant, I. and N.K. Smith. 1933. *Critique of Pure Reason*. London: Macmillan and Co.

Mizrahi, M. 2015. Ought, Can, and Presupposition: An Experimental Study. *Methode-Analytic Perspectives* 4: 232–243.

Moore, G.E. 1922. *Philosophical Studies*. New York: Harcourt, Brace, and Co.

Nichols, S. and J. Knobe. 2007. Moral responsibility and determinism: the cognitive science of folk intuitions. *Nous* 41: 663–85.

Parfit, D. 1984. *Reasons and Persons*. Oxford: OUP.

Pinillos, A., N. Smith, S. Nair, P. Marchetto and C. Mun., 2011. Philosophy's new challenge: experiments and intentional action. *Mind and Language* 26: 115–39.

Sidgwick, H. 1884. *The Methods of Ethics*. London: Macmillan and Co.

Sinnott-Armstrong, W. 1984. 'Ought' conversationally implies 'can'. *The Philosophical Review* 93: 249–61.

Sinnott-Armstrong, W. 1985. 'Ought to have' and 'could have'. *Analysis* 45: 44–48.

Stich, S. and J.M. Weinberg. 2001. Jackson's empirical assumptions. *Philosophy and Phenomenological Research* 62: 637–43.

Stocker, M. 1971. 'Ought' and 'can'. *Australian Journal of Philosophy* 49: 303–16.

Streumer, B. 2003. Does 'ought' conversationally imply 'can'? *European Journal of Philosophy* 11: 219–28.

Streumer, B. 2007. Reasons and impossibility. *Philosophical Studies* 136: 351–84.

Vranas, P. 2007. I ought, therefore I can. *Philosophical Studies* 136: 167–216.

White, A. 1975. *Modal Thinking*. New York: Cornell University Press.

White, M. 1979. Ought and Can. *In The Idea of Freedom*. Oxford: OUP.

Williamson, T. 2010. *The Philosophy of Philosophy*. New York: Blackwell.

Zimmerman, M. 1996. *The Concept of Moral Obligation*. Cambridge: Cambridge University Press.

Abstract

Most philosophers assume that 'ought' implies 'can', and most of them hold that this principle is true not only universally but also analytically or conceptually. Some skeptics deny this principle, although they often admit some related one. In this article, we show how new empirical evidence bolsters the skeptics' arguments. We then defend the skeptical view against some objections to the empirical evidence and to its effect on the traditional principle. In light of the new evidence, we conclude that philosophers should stop unconditionally assuming that 'ought' implies 'can'.