Recent Experimental Work on ‘Ought’ Implies ‘Can’

In press at Philosophy Compass

Jen Semler & Paul Henne

Abstract: While philosophers generally accept some version of the principle ‘ought’ implies ‘can’, recent work in experimental philosophy and cognitive science provides evidence against a presupposition or a conceptual entailment from ‘ought’ to ‘can’. Here, we review some of this evidence, its effect on particular formulations of the principle, and future directions for cognitive scientists and philosophers.

1. Introduction

Philosophers often adhere to some version of OIC:

(OIC) For any agent, A, and act, X, if A ought to X, then A can X.

Kant (1998; cf. Stern 2004; Ranganathan, 2010; Kohl, 2015), Sidgwick (1907), Moore (1922), Parfit (1984), Zimmerman (1996), Vranas (2007), and many others endorse this principle. Some philosophers provide arguments for (Streumer, 2007; 2010; 2018) and against (Brownlee, 2010; Heuer, 2010; Henne & Semler et al., 2018) OIC. Many philosophers, however, simply assume the principle.

Recently, philosophers and cognitive scientists have collected experimental data on people’s judgments about what agents ought to do, and the results of these studies provide evidence against this well-established principle. In this article, we evaluate the empirical evidence against OIC. In section 2, we introduce and review evidence against the claims that OIC is a presupposition or a conceptual entailment. In section 3, we respond to objections that suggest that this empirical evidence does not investigate the relevant philosophical interpretation of ‘ought’. In section 4, we briefly review some evidence for alternative theories about the relationship between ‘ought’ and ‘can’. In section 5, we reply to the claim that empirical evidence is useless in evaluating OIC. Finally, in section 6, we present directions for future research.

2. Against OIC

Philosophers who accept OIC often claim that that ‘ought’ presupposes or conceptually entails ‘can’. In this section, we review empirical work that challenges these positions.

2.1 *OIC is not a Presupposition*

Some philosophers argue that ‘ought’ presupposes ‘can’ (Hare, 1965). On this presupposition view, the statement “A ought to X” presupposes “A can X” in the same way that the statement “the King of France is wise” presupposes “there is a King of France” (Hare, 1965: 54). This relationship is supposed to exclude the possibility that someone ought to do something that he cannot do (Mizrahi, 2015a: 232). As such, this theory yields testable hypotheses: if ‘ought’ presupposes ‘can’, then (1) people’s judgments that “A ought to X” should be correlated with their judgments that “A can X,” and (2) if people judge that “A ought to X,” they should not also judge that “A cannot X.”

Researchers have recently tested whether people’s judgments accord with the presupposition view. In Mizrahi’s study (2015a), participants were given one of four vignettes describing an agent (Professor Smith or Nancy) who promises to meet a student (Sid). The agent becomes unable to fulfill the promise because he either forgets about the meeting or is locked in a room. In one set of vignettes, participants were then asked to what extent they agreed or disagreed with, for instance, the following statements: “Professor Smith *can* keep his office appointment with Sid” and “Professor Smith *ought to* keep his office appointment with Sid.” Generally, participants judged both that Professor Smith ought to keep his appointment and that Professor Smith cannot keep his appointment. Because (1) people’s judgments that the agent ought to keep his promise were not correlated with their judgments about the agent’s ability in any of the vignettes and (2) judgments that the agent ought to act were made in tandem with judgments that the agent cannot act, the study provides evidence against the claim that ‘ought’ presupposes ‘can’ (Mizrahi, 2015a; Kissinger-Knox, Aragon, & Mizrahi, 2017).

2.2 *OIC is Not a Conceptual Entailment*

Other philosophers hold that “A ought to X” conceptually entails “A can X.” On this view, “A can X” is defined such that A has both the ability and the opportunity to X (Vranas, 2007). Suppose we claim,

“John ought to drive Phil to the airport.” On the conceptual entailment view, this claim necessarily implies that John has the required physical and mental ability plus the opportunity to drive Phil to the airport (Vranas, 2007). If ‘ought’ conceptually entails ‘can’, then the contrapositive must be a conceptual entailment as well: ‘cannot’ must entail ‘not-ought.’ Henne and colleagues (2016) note that if this is the case, people should reject the claim that an agent ought to act when the agent is unable to do so. Just as people would readily deny that “A is a bachelor” upon learning that “A is married,” so too should they reject “A ought to X” upon learning that “A cannot X.”

New empirical evidence, however, challenges this conceptual entailment view. Buckwalter and Turri (2015), for instance, examined the relationship between ‘ought’ judgments and inability in various scenarios. In one experiment, participants were given a vignette where an agent had an obligation but was unable to fulfill it. The researchers varied whether the inability was physical (car accident) or psychological (clinical depression) in nature. Surprisingly, in both cases, when people agreed that the agent was unable to act, they also agreed that the agent ought to act. Buckwalter and Turri replicated this result, finding the same pattern regardless of the inability type (Experiments 1, 2, and 3), the phrasing of the moral prompt (Experiment 2), the perspective of the story (Experiment 3), the consequence of the inaction (Experiment 4), the duration of the inability (Experiment 5), and the scope of the inability (Experiment 6). Overall, these studies provide evidence that it is not the case that ‘ought’ conceptually entails ‘can’.

Further studies have explored the conceptual entailment view by looking at the correlation between people’s judgments of ‘ought’ and ‘can’ (Henne et al., 2016; Chituc et al., 2016). If the conceptual entailment view were true, then judgments that “A ought to X” would be strongly correlated with judgments that “A can X.” In one experiment, researchers explored this possibility. Participants read a vignette in which an agent promises to meet his boss on a train, but the agent’s car breaks down, and so he is unable to keep his promise (Chituc et al., 2016: Experiment 2). Critically, there was no evidence of a correlation between participants’ judgments that the agent ought to keep his promise and their judgments that the agent (at that time) can keep his promise. Given these consistent results, philosophers should not assume OIC as a conceptual entailment.

2.3 Explaining the ‘Ought’-‘Can’ Relationship

The above experiments provide reason to believe that it is not the case that ‘ought’ implies ‘can’. If it is true that ‘ought’ does not imply ‘can’, we must still answer the question of why ‘ought’ so often *seems* to imply ‘can’—we must explain why philosophers take OIC to be so intuitive.

While there may be alternative explanations for OIC’s intuitiveness, some have proposed that the blameworthiness of the agent might influence the relationship between ‘ought’ and ‘can’ (Sinnott-Armstrong, 1984; Henne et al. 2016; Chituc et al., 2016). This view has some empirical support. In one experiment, Chituc and colleagues (2016) gave participants a vignette about an agent who is unable to keep his promise to meet his friend because either his car breaks down (low-blame) or he decides to stay home instead (high-blame). Participants were then asked to what extent they agreed with statements that the agent ought to keep his promise, can keep his promise, and is to blame for failing to keep his promise. In all cases, people’s judgments that the agent ought to keep his promise were not correlated with their judgments that the agent can keep his promise. There was, however, a correlation between the extent to which people agreed that the agent ought to keep his promise and the extent to which people agreed that the agent was blameworthy for failing to keep his promise—even in the low-blame cases (Chituc et al., 2016).

These findings can help us understand why OIC is so intuitive to some philosophers. In low-blame cases, a desire to exculpate, rather than the inability itself, seems motivate the retraction of ‘ought’ judgments (Henne et al., 2016; Chituc et al., 2016). The phenomenon of *excuse validation* might help explain this relationship (Henne et al., 2016; Chituc et al., 2016). Generally, excuse validation describes a phenomenon where people deny that an agent violated a rule when the agent is blameless for the rule violation (Turri and Blouw, 2015). As this applies to OIC, people want to excuse blameless agents (in low-blame cases), so people retract their judgments of the agent’s obligation. Hence, it is not inability, but instead the fact that the agent is not blameworthy for the inability, that prompts the denial of the agent’s obligation and the upholding of OIC (Henne et al., 2016; Chituc et al., 2016).

The excuse validation account suggests that blame, rather than inability, drives people’s judgments about what an agent ought to do. This view explains why OIC seems intuitive to many philosophers. More

specifically, the relationship between judgments of what an agent ought to do and judgments of the agent's blameworthiness better explains moral judgments than the relationship between judgments of what an agent ought to do and judgments of what the agent can do. While this explanation seems plausible, there may be alternatives, and exploring these relationships would be a fruitful area for future research.

3. Objections to Empirical Results

Despite these findings and the explanations that they offer, some philosophers continue to accept OIC and reject the recent empirical results for various reasons. For instance, some philosophers claim that the experimental studies fail to capture the meaning of 'ought' that interests philosophers. More specifically, some argue that the participants in the studies might not actually be violating OIC because they are using a different notion of 'ought' than that which is philosophically relevant. In this section, we respond to these objections.

3.1 Nonmoral 'Oughts'

Some philosophers have criticized the empirical OIC research for failing to investigate moral, as opposed to nonmoral, obligations. Kurthy and Lawford-Smith (2015), for instance, criticized Mizrahi's (2015a) experimental design on the grounds that the use of 'ought' was ambiguous; participants may have falsely interpreted 'ought' as optative, describing a wish rather than a moral obligation. For example, when participants judged that the agent ought to keep his promise, their judgments may have indicated that the agent would keep his promise in an ideal world.

In response to this criticism, Mizrahi (2015b) replicated the findings of the original (2015a) study and—in order to avoid the potential confound—attempted to elicit an all-things-considered moral judgment in participants. In order to make participants more likely to interpret the 'ought' as moral, Mizrahi changed the wording of the question specifically to target keeping the promise (a moral act) rather than attending the meeting (a nonmoral act). With this change, people still judged that the agent ought to keep his promise despite participants' knowledge of the agent's inability—even when they were prompted to respond to an all-things-considered moral 'ought' (Mizrahi, 2015b).

There may be other potential confluences of distinct ‘oughts’, but recent studies do not support these claims. Buckwalter and Turri (2015), for example, tested whether participants could distinguish between moral and legal obligations. It is possible that a legal ‘ought’ (what an agent ought to do in virtue of the law) may imply ‘can’ where the moral ‘ought’ (what an agent ought to do in virtue of morality) does not, or vice versa. Hence, it is important to know whether people conflate these concepts. In one experiment, participants read vignettes in which a character was either able or unable to save a drowning child (Buckwalter and Turri, 2015: Experiment 7). Some participants were asked whether the agent was morally obligated to save the child, and others were asked whether the agent was legally obligated to save the child. While participants made judgments inconsistent with OIC in the moral domain (by judging that the agent ought to but cannot act), their judgments were consistent with OIC in the legal domain (by judging that the agent was neither able to nor legally obligated to act). These results indicate that people distinguish different types of ‘oughts’ and that they treat different formulations of OIC differently in distinct domains. People seem to distinguish moral obligations from other, similar types of obligations, and there is no evidence that they are conflating distinct concepts in the experiments. Hence, if philosophers continue to think that participants in these studies are failing to respond to a moral ‘ought’, they will have to provide further evidence for this claim.

3.2 *Multiple Meanings of ‘Ought’*

There may still be some ambiguity regarding the meaning of moral ‘oughts’, and this potential confound could affect the interpretation of the empirical results. So, when people make judgments that are inconsistent with OIC, some philosophers may wonder if participants are using a sense of ‘ought’ other than that which is philosophically relevant. In recent work, Turri (2017) considers just this; specifically, he explores the idea that ‘ought’ can have a descriptive purpose (moral obligation) or an encouragement purpose (promoting behavior) (see Sinnott-Armstrong, 1984; 1985). In one study, he presented participants with vignettes in which Alex promises to deliver a package for Julia, and Julia says that Alex ought to deliver the package. (Turri, 2017: Experiment 1). Some participants were given a scenario in which Alex is able to keep his promise, while other participants were given a scenario in which Alex is unable to keep

his promise. Participants then rated the extent to which Julia's statement describes Alex's moral obligation and the extent to which Julia's statement encourages Alex to act. The results indicate that descriptive judgments were the same in all cases: regardless of whether Alex was able to deliver the package, people judged that Alex's sister described a moral obligation. Encouragement judgments, on the other hand, varied depending on the agent's ability: when an action was possible—but not when it was impossible—people rated the statements about what Alex ought to do as promoting behavior.

If people's judgments about what an agent ought to do function differently in distinct contexts, it is important to evaluate how different senses of 'ought' are related to ability and how these results affect the overall findings against OIC. Using a timeline activity, Turri (2017: Experiment 2) tested how people's encouragement and descriptive judgments changed before and after an inability made an action impossible. Turri found that when people rated 'ought' as descriptive, people also judged that the obligation extended temporally beyond the agent's ability: even after the agent became unable to perform the act, people still judged that the agent had an obligation to act. When people rated 'ought' as encouragement, however, inability was a relevant constraint: participants judged that once it became impossible for the agent to perform an act, it was no longer worthwhile for the agent to be encouraged to perform it. These findings suggest that there are at least two senses of 'ought' used in people's ordinary moral judgments and that judgments about what an agent ought to do are constrained by ability only if people interpret 'ought' as encouraging a specific action. Ultimately, Turri's studies again confirm the finding that people's judgments are inconsistent with OIC—even when they are carefully prompted to consider the moral-obligation sense of 'ought' that so often concerns philosophers.

3.3 *'Ought' as 'Ought to Have'*

There are other ways that participants might have misinterpreted 'ought' which could affect the conclusions drawn from these studies. For example, Hannon (2017) argues, specifically in the experiments performed by Chituc and colleagues (2016), that participants could have wrongly interpreted 'ought' as 'ought to have'. Hence, when participants judged that an agent 'ought' to perform an act, they responded affirmatively in the sense that that the agent 'ought to have' performed the act before the agent was unable

to perform the act. If this is the case, the perceived inability of an agent might be relevant to the judgment that “A still ought to X,” but it might be irrelevant to the judgment that “A ought to have done X when A was able to X.” Hannon tested this hypothesis using the same high-blame scenario as Chituc and colleagues (2016) in which an agent intentionally stays at home so that he cannot fulfill his promise to meet a friend. Hannon added a follow-up question to investigate whether participants’ judgments were meant to signal that the agent ‘still ought to’ perform the impossible action or that they ‘ought to have’ left the house earlier to fulfill the promise. The results replicated those from Chituc and colleagues’ high-blame scenario. However, in the follow-up question, 58% of participants who initially made judgments inconsistent with OIC clarified that they were making an ought-to-have evaluation rather than a still-ought-to evaluation.

While this study raises some doubts for the tested interpretation of ‘ought’, there are two significant issues with it. First, because participants were asked to choose one option (‘still ought’, ‘ought to have’, or neither), it is unclear whether some participants who selected the ‘ought to have’ interpretation also believed that the agent still had an obligation after the agent was no longer able to keep the promise (see also Sinnott-Armstrong, 1985). Second, in response to the follow-up question, a significant percentage of participants (34%) held that the agent “still ought” to perform the act that he is unable to perform (Hannon, 2017). Defenders of OIC must explain why some people still make both “still ought” and inability judgments in certain situations without appeal to a mistaken interpretation of ‘ought’.

3.4 *Time of Inability and Obligation*

There may also be a concern about study participants considering the moral obligation at a time other than which is specified by the experimenters (Cohen, 2018; Streumer, 2018). Leben (2018) attacks earlier studies on these grounds and argues that temporal vagueness in the study designs led participants to make judgments inconsistent with OIC. To show this, Leben conducted an experiment in which Buckwalter and Turri’s (2015) vignettes were represented in a timeline diagram. For example, Leben placed “Walter promises Brown,” “car accident,” and “Brown not picked up” in a visual sequential order. He put an arrow in the diagram pointing to the moment in the story where participants should evaluate ‘ought’ and ‘can’ questions. So, the judgments were temporally located either before or after the inability-causing accident

occurred, rather than at the time of the accident. Notably, participants agreed with the ‘ought’ statement at a significantly higher rate when it referred to a time *before* the inability compared to when it referred to a time *after* the inability. Leben concludes that ability influences ‘ought’ judgments when time is explicitly stated (Leben, 2018).

There are, nonetheless, some potential issues with this study. First, there seem to be major demand characteristics at work. By presenting the timeline visually, as opposed to using vignettes, it may be more obvious to participants that experimenters were testing for differences in ‘ought’ judgments at different times. This manipulation could influence participants’ behavior and encourage them to make judgments more consistent with OIC simply because they think the experimenters want them to respond accordingly. Second, in Leben’s first study, the mean ‘ought’ rating at the time after the inability was 2.58 on a 0-6 scale (higher ratings indicate a more affirmative ‘ought’ judgment). Leben acknowledges that this finding casts doubt on a strong notion of OIC, as a conceptual entailment would require ‘ought’ ratings closer to 0 when an inability is present (160). Leben, nonetheless, explains this discrepancy by proposing that the term ‘ought’ has more than one meaning: it can indicate a prior obligation or an all-things-considered moral judgment. And these distinct meanings, according to Leben’s view, might influence how people evaluate OIC in many of the studies.

Hence, Leben examines how the phrasing of the responses from which participants select might lead to a mistaken understanding of the time at which ‘ought’ is supposed to be evaluated. Specifically, he suggests that if ‘ought’ is located before ‘can’ in a sentence, people might interpret ‘ought’ as indicating a previous obligation instead of a current obligation. For example, on this view, the sentence “A ought to X, but A cannot X” would be interpreted as “A had a previous obligation to X, but A is currently unable to X” instead of “A has a current obligation to X, but A is currently unable to X”. Leben again used Buckwalter and Turri’s (2015) materials with both the original and the reordered statements (for example, original: “Walter is obligated to pick up Brown at the airport, but Walter is not physically able to do so” and reordered: “Walter is not physically able to pick Brown up at the airport, but he is still obligated to do so”). The results of some of the scenarios show that when the original ordering was reversed, fewer participants

chose the inability-obligation response (Leben, 2018). These findings are consistent with those of Kurthy et al. (2017), who similarly altered Buckwalter and Turri's (2015) designs by changing the wording of the response options and found that people made fewer OIC violations.

There are a few reservations to keep in mind about these results. First, the experiment shows that there is an ordering effect on people's judgments about whether an agent ought to act. The results, however, do not establish what specifically causes the difference in judgments or which ordering is the right set of judgments to evaluate in relation to OIC. The reordering is assumed to remove a bias or vagueness in the time participants are asked about, but this claim is not substantiated by the results of these studies, which show only a difference in judgments. Second, the use of the words "but" and "still" in the prompts may also influence the participants' choices. The inclusion of these connectives creates potential demand characteristics, as participants might think that the experimenters are prompting them to answer in a way consistent with OIC. Thus, while these results yield a strong start to understanding temporality and OIC judgments, more work is required for generalizability.

4. Alternative Principles

Given these consistent findings about the lack of a relationship between 'ought' and 'can', some philosophers who reject OIC have suggested alternative principles that better explain the role of ability in moral judgments. We explore these views here.

4.1 'Oughts' and 'Cans' on Multiple Levels

In light of the empirical results discussed thus far, some philosophers have offered alternative interpretations of OIC. Willemsen and Wiegmann (2017) suggest that 'ought' is used pragmatically, and it is used to convey two senses of 'ought': 'ought' in the sense of a moral obligation and 'ought' in the sense of an imperative. In one of their experiments, participants were given one of two high-blame vignettes, in which an agent knowingly chooses actions that will lead to him breaking a promise (Experiment 2). Participants were randomly assigned to one of three conditions and asked whether they agree or disagree

with the statement that the agent ‘ought to’ (original condition), ‘has to’ (imperative condition), or ‘has the moral obligation to’ (obligation condition) carry out his promise.

Most participants (68% and 81%) agreed with the obligation condition, while few agreed with the imperative condition (29% and 32%). In the original condition, agreement rates fell between those of the other two conditions (45% and 74%). These findings are consistent with Turri’s (2017) results regarding the function of ‘ought’ as descriptive and encouragement. Willemsen and Wiegmann took these results as evidence that people use ‘ought’ to express different meanings. From these findings, the authors propose that OIC may be upheld if it operates on two levels, where ‘ought’ as an obligation implies one sense of ‘can’, while ‘ought’ as an imperative implies a different sense of ‘can’. The experimenters did not test interpretations of ‘can’ judgments—they proposed different interpretations of ability to fit with their two-level theory. Further research is necessary to determine the nature of these relationships and how they provide evidence for or against OIC.

4.2 *Blame Implies Can*

Other researchers have suggested that blame implies can (BIC) may be a more suitable principle because, while they have failed to find a strong association between ‘ought’ and ability, studies have found correlations between judgments about blame and ability (Buckwalter & Turri, 2015; Henne et al., 2016; Chituc et al., 2016). BIC holds that if an agent is blameworthy for an action, then the agent must be able to perform that action. In its contrapositive form, this principle indicates that if an agent cannot perform an action, then the agent is not blameworthy for its failure.

Buckwalter and Turri (2015) actually compared BIC and OIC in an experiment. Participants read a vignette about an agent who was physically unable to keep his promise because he got in a car accident on his way to pick up his friend. In the BIC condition, participants were asked whether the agent was blameworthy for not keeping the promise; in the OIC condition, participants were asked whether the agent ought to keep the promise. In the BIC condition, 97% of participants responded that the agent was not blameworthy and unable to act, but in the OIC condition, 85% responded that the agent ought to keep his

promise but was unable to act. These results provide preliminary evidence that BIC might provide a better description of moral reasoning than OIC, thus the principle should be explored in future research.

5. Can Experiments Tell Us Anything About OIC?

In a standard critique of any result in experimental philosophy, some philosophers doubt that empirical results have value in determining the status of OIC. Many appeal to the “expertise defense” by claiming that non-philosophers’ moral judgments are irrelevant to the truth of moral principles like OIC (see Mizrahi, 2015c: 10; Cappelen, 2012). These philosophers often argue that because experts (i.e., philosophers) understand moral theory better than non-philosophers, philosophers’ moral judgments about hypothetical cases have more epistemic weight. However, evidence has shown that philosophers are, like non-philosophers, subject to order effects (Schwitzgebel & Cushman, 2012; 2015), framing effects (Tobia et al., 2013a), and cleanliness effects (Tobia et al., 2013b)—just to name a few. Even reflectiveness, often thought to be at the core of philosophical thinking, does not seem to change judgments about philosophical thought experiments (Colaço et al., 2018; see also, Machery, 2017). Because philosophers are subject to the same biases as non-experts, the expertise defense is not very convincing.

Other critics argue that because ordinary judgments are flawed, they do not affect OIC. Kissinger-Knox et al. (2017) claim that empirical studies can, at most, show that OIC is unintuitive. Yet, they argue, empirical studies cannot refute OIC. Their experiment shows an actor-observer bias, whereby people’s judgments about ‘ought’ and ability differ depending on whether the person making the judgment is the actor or the observer. They conclude that because OIC judgments can be distorted by factors that should be irrelevant, the judgments are unreliable indicators of the truth of the principle. However, we should not abandon OIC empirical studies in light of this bias (and other cognitive biases). Additionally, these results indicate that judgments about OIC can be biased in specific ways, but this does not mean that all judgments about OIC are flawed.

6. Future Work

The existing studies on OIC leave many avenues open for future research. The empirical evidence has shown that OIC is not obviously true. But there is a lot of work to be done. First, future research should continue to examine the relationship between blame, moral obligation, and ability—specifically how blame and exculpation relate to the distinct functions of ‘ought’—and how these effects relate to recent work on counterfactual thinking. Second, the basis for empirical refutations of OIC is that people make judgments that violate the principle. Future cross-cultural studies can shed light on the universality of OIC by determining how judgments about what agents ought to do are related to judgments about ability in other cultures. Lastly, future studies should expand the types of moral judgments that are tested. Previous studies have largely focused on promises. Researchers should investigate whether OIC applies to other types of moral obligations such as honesty, fairness, and loyalty.

Additionally, while previous work has been concerned with whether agents ought to act when they cannot, future studies should explore OIC beyond the realm of action. Recent work reveals a similar pattern of results in different domains. People judge agents as having some duty, or intellectual responsibility, to hold certain beliefs based on evidence and objective truth even if they are unable to do so (Buckwalter and Turri, 2018). Future studies should explore whether there is a moral element to these judgments about intellectual responsibility. In other words, researchers should test whether people judge that agents morally ‘ought to believe’ or ‘ought to know’ something that they are unable to believe or know. Similarly, future studies should explore how claims about how an agent ‘ought to be’ or ‘ought to feel’ are influenced by inability. These future projects suggest that there is a surplus of empirical work to be done on OIC.

7. Conclusion

Philosophers have been discussing OIC roughly since Kant. Yet most tend to assume the principle without much investigation. Only recently have experimental philosophers and cognitive scientists started looking at people’s judgments about these concepts. The evidence so far suggests that people do not consistently make judgments in accordance with OIC. At the very least, this means that philosophers should

stop assuming the principle. The burden of proof has shifted to defenders of OIC to reconcile their views with the empirical evidence and start providing new arguments for OIC or a new principle altogether.

Acknowledgements

We would like to thank Fiona Xin, Irene Zhou, Walter Sinnott-Armstrong, Felipe De Brigard, and Duke's MAD Lab for comments on this manuscript.

References

- Buckwalter, W., & Turri, J. (2015). Inability and obligation in moral judgment. *PloS one*, 10(8), e0136589.
- Buckwalter, W., & Turri, J. (2018) Inability and obligation in intellectual evaluation. *Episteme*, 1-23. 10.1017/epi.2018.49.
- Buckwalter, W. (2017). Ability, responsibility, and global justice. *Journal of Indian Council of Philosophical Research*, 34(3), 577-590.
- Brownlee, K. (2010). Moral aspirations and ideals. *Utilitas*, 22(3), 241-157.
- Cappelen, H. (2012). *Philosophy without intuitions*. Oxford: Oxford University Press.
- Chituc, V., Henne, P., Sinnott-Armstrong, W., & De Brigard, F. (2016). Blame, not ability, impacts moral “ought” judgments for impossible actions: Toward an empirical refutation of “ought” implies “can”. *Cognition*, 150, 20-25.
- Cohen, Y. (2018). An analysis of recent empirical data on ‘ought’ implies ‘can’. *Philosophia*, 46(1), 57-67.
- Colaço, D., Kneer, M., Alexander, J., & Machery, E. (2018). On second thought: A refutation of the reflection defense. 10.13140/RG.2.2.34481.68967.
- Hannon, M. (2017). Intuitions, reflective judgments, and experimental philosophy. *Synthese*, 1-22.
- Hare, R. M. (1965). *Freedom and reason*. Oxford: Oxford University Press.
- Henne, P., Chituc, V., De Brigard, F., & Sinnott-Armstrong, W. (2016). An empirical refutation of ‘ought’ implies ‘can’. *Analysis*, 76(3), 283-290.

- Henne, P., Semler, J., Chituc, V., De Brigard, F., & Sinnott-Armstrong, W. (2018). Against some recent arguments for 'ought' implies 'can': Reasons, deliberation, trying, and furniture. *Philosophia*. <https://doi.org/10.1007/s11406-017-9944-7>.
- Heuer, U. (2010). Reasons and impossibility. *Philosophical Studies*, 147, 235-246.
- Kant, I. (1998). *Critique of Pure Reason* (translated and edited by Paul Guyer & Allen W. Wood).
- Kissinger-Knox, A., Aragon, P., & Mizrahi, M. (2017). "Ought implies can," framing effects, and "empirical refutations". *Philosophia*, 46(1), 165-182.
- Kohl, M. (2015). Kant and 'ought implies can'. *The Philosophical Quarterly*, 65(261), 690-710.
- Kurthy, M., & Lawford-Smith, H. (2015). A brief note on the ambiguity of 'ought'. Reply to Moti Mizrahi's 'Ought, can and presupposition: An experimental study'. *Methode*, 4(6), 244-249.
- Kurthy, M., Lawford-Smith, H., & Sousa, P. (2017). Does ought imply can?. *PloS one*, 12(4), e0175206.
- Leben, D. (2018). In defense of 'ought implies can'. *Oxford Studies in Experimental Philosophy*, 2, 151.
- Machery, E. (2017). *Philosophy within its proper bounds*. Oxford: Oxford University Press.
- Mizrahi, M. (2015a). Ought, can, and presupposition: An experimental study. *Methode*, 4, 232-243.
- Mizrahi, M. (2015b). Ought, can, and presupposition: A reply to Kurthy and Lawford-Smith. *Methode*, 4(6), 250-256.
- Mizrahi, M. (2015c). Three arguments against the expertise defense. *Metaphilosophy*, 46(1), 52-64.
- Moore, G. E. (1922). *Philosophical Studies*. New York: Harcourt, Brace, and Co.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Oxford University Press.
- Ranganathan, S. (2010). Does Kant hold that ought implies can?. *Grounding Morality: Freedom, Knowledge and the Plurality of Cultures*, 60.
- Schwitzgebel, E., & Cushman, F. (2012). Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind & Language*, 27, 135-153.
- Schwitzgebel, E., & Cushman, F. (2015). Philosophers' biased judgments persist despite training, expertise and reflection. *Cognition*, 141, 127-137.

- Sinnott-Armstrong, W. (1984). 'Ought' conversationally implies 'can'. *The Philosophical Review*, 93(2), 249-261.
- Sinnott-Armstrong, W. (1985). 'Ought to have' and 'could have'. *Analysis*, 45(1), 44-48.
- Sidgwick, H. (1907). *The methods of ethics*. Hackett Publishing.
- Stern, R. (2004). Does 'ought' imply 'can'? And did Kant think it does?. *Utilitas*, 16(1), 42-61.
- Streumer, B. (2007). Reasons and Impossibility. *Philosophical Studies*, 136, 351-384.
- Streumer, B. (2010). Reasons, impossibility, and efficient steps: Reply to Heuer. *Philosophical Studies*, 151, 79-86.
- Streumer, B. (2018). I ought to reply, so I can. *Philosophia*. <https://doi.org/10.1007/s11406-018-0042-2>
- Tobia, K., Buckwalter, W., and Stich, S. (2013a). Moral intuitions: Are philosophers experts? *Philosophical Psychology*, 26, 629-638.
- Tobia, K., Chapman, G. B., and Stich, S. (2013b). Cleanliness is next to morality, even for philosophers. *Journal of Consciousness Studies*, 20, 195-204.
- Turri, J. (2017). How "ought" exceeds but implies "can": Description and encouragement in moral judgment. *Cognition*, 168, 267-275.
- Turri, J., & Blouw, P. (2015). Excuse validation: A study in rule-breaking. *Philosophical Studies*, 172(3), 615-634.
- Vranas, P. B. (2007). I ought, therefore I can. *Philosophical Studies*, 136(2), 167-216.
- Willemsen, P., & Wiegmann, A. (2017). I must although I can't!? Suggestions for a two-level theory of 'ought implies can'. <https://doi.org/10.31234/osf.io/hyq9u>.
- Zimmerman, M. (1996). *The Concept of Moral Obligation*. Cambridge: Cambridge University Press.