

An Analysis of Recent Empirical Data on ‘Ought’ Implies ‘Can’

Yishai Cohen
Published in *Philosophia*

Abstract Recent experimental studies dispute the position that commonsense morality accepts ‘Ought’ Implies ‘Can’ (OIC), the view that, necessarily, if an agent ought to perform some action, then she can perform that action. This paper considers and supports explanations for the results of these studies on the hypothesis that OIC is intuitive and true.

Keywords Ability, Blame, Blame validation, Experimental philosophy, Obligation, Ought implies can

1.

When non-philosophers appear to disagree with a philosopher’s theory, the philosopher should provide a plausible explanation for why this apparent disagreement exists on the hypothesis that the philosopher’s theory is true. A number of recent experimental studies explore the extent to which commonsense morality accepts ‘Ought’ Implies ‘Can’ (OIC), the principle that every possible objective,¹ all-things considered moral obligation is accompanied by an ability (and opportunity) to fulfill that obligation.² This principle may also be formulated as follows:

(OIC) Necessarily, if at time t an agent A ought to ϕ , then at t A can ϕ .

According to the experimental studies of Wesley Buckwalter and John Turri (2015), Moti Mizrahi (2015a; 2015b), and Vladimir Chituc, Paul Henne, Walter Sinnott-Armstrong, and Felipe De Brigard (2016) commonsense morality does not accept OIC. By contrast, the studies of Miklos Kurthy, Holly Lawford-Smith, and Paulo Sousa (2017) and Michael Hannon (Forthcoming) support the opposite conclusion. This paper defends and expands upon the position of Kurthy *et al.* and Hannon by scrutinizing the evidential import of this empirical data upon OIC, starting with the studies by Buckwalter and Turri.

2.

¹ An objective obligation is one that depends upon the normatively relevant facts, regardless of whether the agent has access to those facts. For instance, suppose that, unbeknownst to Robert, pressing the elevator button will result in a deadly explosion. All other things being equal, Robert has an objective obligation to refrain from pressing the button, even if he is not aware of the consequences of pressing the button. By contrast, a subjective obligation is one that depends upon the evidence available to the agent. For instance, if Robert has a justified false belief that pressing the elevator button will result in a deadly explosion, then (all other things being equal) Robert has a subjective obligation to refrain from pressing the button. Since it is possible for an agent to form a justified false belief that they can ϕ , it is possible for an agent to have a subjective obligation that they cannot in fact fulfill. For this reason, OIC is formulated in terms of an objective obligation (cf. Vranas (2007) and Graham (2011)).

² Other formulations of OIC consider ‘ought’ to refer to an objective, *pro tanto* obligation, rather than an all-things-considered obligation (cf. Vranas 2007). This difference is peripheral to the discussion at hand.

In Buckwalter and Turri's (2015: 3-5) Experiment 1 participants are provided with one of two vignettes:

Walter promised that he would pick up Brown from the airport. But on the day of Brown's flight, Walter is [in a serious car accident/suffering from clinical depression]. As a result, Walter is not [physically/psychologically] able to pick up Brown at the airport.

Participants must then choose one of the following four options that best applies to the vignette:

1. Walter is obligated to pick up Brown at the airport, but Walter is not physically able to do so.
2. Walter is not obligated to pick up Brown at the airport, and Walter is not physically able to do so.
3. Walter is obligated to pick up Brown at the airport, and Walter is physically able to do so.
4. Walter is not obligated to pick up Brown at the airport, but Walter is physically able to do so.

Since 80% of participants selected the first option with respect to the 'car accident' vignette, Buckwalter and Turri take this result to serve as a counterexample to OIC.

Kurthy *et al.* (2017) hypothesize that participants who chose the first option are not in fact denying OIC for the following reasons. The instruction to "choose the option that best applies" may suggest to participants that they are being tested on their ability to comprehend the story, rather than being tested on their views concerning the conditions under which an agent possesses an obligation. Moreover, the first option best captures the temporal dimension of the story, viz. that Walter *was* obligated to do something that he *subsequently* is unable to do. Since this claim is consistent with OIC, participants who selected the first option in order to capture this temporal dimension of the story are not committed to denying OIC. Call this the *temporal dimension* hypothesis.

Kurthy *et al.* (2017: 4-6) support this hypothesis by conducting a similar experiment with the same 'car accident' vignette but with slightly different answers (Study 1). In the responses available to the participants, the connectives "and" and "but" are replaced with "even if" and "because" in order to highlight OIC to the participants:

- A. Walter is obligated to pick up Brown at the airport, even if Walter is not physically able to do so.
- B. Walter is not obligated to pick up Brown at the airport, because Walter is not physically able to do so.

In contrast to the study by Buckwalter *et al.*, 88% of participants in this study selected (B) which in turn supports the *temporal dimension* hypothesis. Kurthy *et al.* (2017: 7-9) further support their hypothesis in a separate study (Study 2) that replicated the original study by Buckwalter and Turri with the original "and" and "but" connectives, but with an important follow-up question. 74.4% of participants who selected the "obligated, but not able" response in Study 2 were subsequently invited to answer the following question:

You chose the option “Walter is obligated to pick up Brown at the airport, but Walter is not physically able to do so.” With this choice, do you mean that Walter is still under the obligation to pick up Brown at the airport after he becomes physically unable to do so? (Yes/No)

The majority of these participants (74%) selected “No”, and then clarified their position in writing. Many of them indicated that they did not intend to affirm that Walter *still* has the obligation once he is physically unable to fulfill it. The *temporal dimension* hypothesis of Kurthy *et al.* thus appears to explain away the apparent anti-OIC intuitions of non-philosophers in the studies of Buckwalter and Turri.

The next section will discuss how the *temporal dimension* hypothesis bears upon Mizrahi’s (2015b) study. But before moving on, this paper introduces some terminology that will be useful for the remainder of our discussion.

An agent A pre-inability-ought to ϕ iff, at a time prior to A’s inability to ϕ , A ought to ϕ . An agent A post-inability-ought to ϕ iff, at a time posterior to A’s inability to ϕ , A ought to ϕ . So according to the *temporal dimension* hypothesis, participants only *appeared* to deny OIC in order to capture the temporal dimension of the vignette, according to which the agent pre-inability-ought to fulfill his obligation, even though he subsequently loses that ability.³

3.

In order to challenge the *temporal dimension* hypothesis, a study must highlight the distinction between the pre-inability-ought and the post-inability-ought notions. Keeping this *desideratum* in mind, let’s turn to Mizrahi’s (2015b: 253) study that appeals to the following vignette:

At the end of class, Sid approaches Professor Smith to ask a question. Professor Smith tells Sid that he has to go teach another class now but promises to meet with Sid during office hours later that day. As it turns out, however, Professor Smith gets locked in his classroom before he is able to make it to his office. After waiting for thirty minutes and realizing that Professor Smith doesn’t show up, Sid leaves the office without meeting Professor Smith.

Each participant is given one of the following statements and then asked to report the extent to which they agree with that statement.

CAN Taking into consideration all the facts about this case, Professor Smith can keep his promise to meet with Sid.

OUGHT Taking into consideration all the facts about this case, Professor Smith ought to keep his promise to meet with Sid.

³ A referee rightly notes that this distinction is not sufficiently fine-grained for cases in which an agent is no longer able to do something but can nevertheless regain this ability (at least under a certain conception of abilities). Since none of the vignettes considered in this paper involve an agent who can regain a lost ability, this distinction is arguably sufficiently fine-grained for the purpose of evaluating the experimental studies under consideration.

BLAME Taking into consideration all the facts about this case, Professor Smith is to blame for failing to keep his promise to meet with Sid.

The available responses range from 1 (= Strongly Disagree) to 5 (=Strongly Agree). The mean ratings for the responses are as follows: CAN: M = 2.03, OUGHT: M = 3.64, BLAME: M = 2.6. Since there is a significant divergence of agreement with respect to all three claims, Mizrahi (2015b: 254) claims that “OIC can no longer be taken as axiomatic. It must be argued for without appealing to intuitions.”

This experiment is intended to address critical commentary by Kurthy and Lawford-Smith (2015) on Mizrahi’s previous (2015a) study, including the concern that participants may not have had in mind an objective, all-things-considered moral ‘ought’ when responding to the vignette. Mizrahi’s (2015b) addresses this concern by incorporating the phrase ‘taking into consideration all the facts about this case’ in the statements to which participants are invited to agree or disagree, and by incorporating into the vignette the moral notion of keeping a promise. Nevertheless, Mizrahi’s study does not obviously satisfy the aforementioned *desideratum*.

The OUGHT claim does not highlight the contrast between the claim that Smith pre-inability-ought to keep his promise and the claim that Smith post-inability-ought to keep his promise. The phrase ‘taking into consideration all the facts about this case’ might suggest to participants that according to OUGHT Smith ought to *have* kept his promise (pre-inability-ought), and it might suggest that according to OUGHT Smith *still* ought to keep his promise once he is locked inside his classroom (post-inability-ought).

The CAN claim that incorporates the same phrase similarly gives rise to an ambiguity: does the phrase mean that Smith can keep his promise *prior* to being locked inside his classroom, or that he can keep his promise once he is locked inside his classroom? These temporal ambiguities are amplified by the fact that ‘taking into consideration all the facts about this case’ suggests to participants that they are being tested on their ability to comprehend the story, rather than being tested on their views concerning the conditions under which an agent possesses an obligation. Like Buckwalter and Turri’s study, participants in this study who seek to display their comprehension of the story may be disposed to affirm that Smith ought to have fulfilled his obligation, even though he cannot do so since he is locked inside his room. So it is not surprising to see a higher rating for OUGHT in comparison to CAN on the hypothesis that OIC is intuitive.

Mizrahi (2015b: 254) also claims that “the higher ratings for the “blame” question compared to the “can” question are not what we would expect if OIC were intuitive.” But it is not clear how these results conflict with OIC either since CAN is temporally ambiguous in a manner that BLAME is not. To illustrate, suppose that Professor Smith is presently blameworthy for not meeting Sid because Smith is presently blameworthy for being locked inside his classroom, and Smith could have avoided being locked inside his classroom, and thus *could have* met Sid prior to being locked inside. Since Smith is presumably not blameworthy *prior* to being locked inside his own classroom (a time at which Smith can still meet Sid), the BLAME claim is not temporally ambiguous in the manner that the CAN claim is ambiguous. As a result, under the interpretation of CAN that says that Smith *could have* met Sid prior to being locked inside, a proponent of OIC can consistently affirm both CAN and BLAME. By contrast, under an interpretation of CAN that says that, at the time at which Smith is locked inside his classroom, he *can* still meet Sid, a proponent of OIC can consistently deny CAN and affirm BLAME. So the

higher ratings for BLAME are also not unexpected if OIC is intuitive. The *temporal dimension* hypothesis thus remains largely unscathed.

4.

The *temporal dimension* hypothesis must also confront the studies of Chituc *et al.* (2016), as well as the elaboration of these studies by Henne *et al.* (2016). In Experiment 1 Chituc *et al.* (2016: 21) present participants with the following two vignettes:

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.

Participants respond to the following question: ‘Do you agree or disagree with the following statement: At eleven forty-five, it is still true that Adams ought to meet Brown at noon.’ The answers are on a scale from -50 (completely disagree) to 0 (neither agree nor disagree) to 50 (completely agree). While 68% of participants denied that Adams ought to keep his promise in *Low Blame*, 60% of participants affirmed that Adams ought to keep his promise in *High Blame*.

The *High Blame* question specifies the time at which it may be true that Adams ought to meet Brown (eleven forty-five), and it contains the word “still” which suggests that Adams’ inability is rendered salient to the reader. This study thus poses a challenge to the *temporal dimension* hypothesis since the *High Blame* question appears to highlight the distinction between the claim that Adams pre-inability-ought to meet Brown at noon and the claim that Adams post-inability-ought to meet Brown. Nevertheless, recent studies by Hannon (Forthcoming) appear to vindicate the *temporal dimension* hypothesis.

Hannon conducted a similar study that incorporates the *High Blame* case, and invited participants to agree or disagree with the *High Blame* question. Like the original study, the majority of participants in Hannon’s study (70%) agreed that at eleven forty-five Adams ought to meet Brown at noon. However, participants were then invited to select one of three claims that best describes their position:

(A) At eleven forty-five, it is still true that Adams ought to meet Brown at noon. It is impossible for him to do this (because it would now take too long to get there on time), but he is obligated to do what is now impossible.

(B) At eleven forty-five, it is true that Adams ought to have met Brown at noon, even though it is now impossible for him to do this (because it would now take too long to get there on time). This is because Adams should have left earlier.

(C) Neither (A) nor (B) adequately describes my opinion (Hannon Forthcoming: 7).

A majority of participants who agreed with the *High Blame* question (58%) selected (B) which says that Adams pre-inability-ought to meet Brown at noon. The second part of Hannon's study thus revealed that the majority of participants did not have anti-OIC judgments.

Hannon (Forthcoming: 6) suggests that the participants who agreed with the *High Blame* question but who selected (B) had a "muddled" intuition in the sense that these participants had an intuition that concerns a different subject matter from the one under consideration (cf. Schroer and Schroer 2013). Specifically, these participants still had the pre-inability-ought notion in mind when responding to the *High Blame* question, even though this question was intended to pick out the post-inability-ought notion. Hannon's study thus appears to corroborate the *temporal dimension* hypothesis of Kurthy *et al.*: participants formed a muddled intuition in order to capture the temporal dimension of the relevant vignette for the sake of displaying their comprehension of the vignette. Moreover, only the second part of Hannon's study highlighted to a sufficient extent the distinction between the pre-inability and post-inability 'oughts', such that participants were more cognitively sensitive of the fact that they were being tested on their views concerning the conditions under which an agent possesses an obligation.

5.

An OIC proponent might seek out further viable hypotheses concerning these studies that corroborate the *temporal dimension* hypothesis. One theory worth considering says that participants are subject to blame validation, a process in which someone who is motivated to blame will exaggerate the volitional control of the actor, lower their standards for blame, or seek information to support their blame attribution (Alicke 2000).⁴ Consequently, such participants have a further motivation for selecting a response that best captures the temporal dimension of the story: by forming a muddled intuition that concerns the pre-inability 'ought', participants successfully seek out a justification for attributing blame to the agent, viz. that the agent *could have* done something that they *should have* done. For example, by forming a muddled intuition that Adams pre-inability-ought to meet Brown at noon, participants claimed to agree with the *High Blame* question in virtue of their disposition to seek out a justification for blaming Adams, viz. that Adams *was* obligated to do something that he could have done, but failed to do it. Call this the blame validation hypothesis.⁵

Chituc et al. consider and reject this hypothesis due to another study, Experiment 2. Participants in this study read the following vignette:

Brown is a CEO of a large company in the economic boom in the middle of the 20th Century. At 2 o'clock, Brown has a meeting in the city to make a significant financial decision that will decide the future of his company. Since so much money is at stake, he asks his trusted personal advisor, Adams, to meet him on the 12 o'clock train. On the

⁴ Alicke's theory may receive support from the hypothesis that affect contributes to people's intuitions (Greene *et al.* 2001; Nichols 2002), and the hypothesis that concrete cases involving wrongdoing are likely to elicit blame-involving emotions (Nichols and Knobe 2007). But this is only a tentative suggestion.

⁵ The 'blame validation hypothesis' is not the hypothesis that people are sometimes subject to the process of blame validation in some experimental study or other. Rather, it's the hypothesis that participants in the experimental studies under consideration are subject to the process of blame validation.

train, he plans to discuss his decision on the ride into the city, where Brown will go straight to his 2 o'clock meeting. Adams promises to meet Brown on the train at noon. It takes Adams thirty minutes to drive to the train station, park, purchase a ticket, and board the train. However, fifteen minutes after leaving at eleven-thirty, Adams car breaks down unexpectedly. Because his car is not working at the time, Adams cannot meet Brown at noon, as promised. Since cell phones have not been invented yet, Adams has no way to contact him (Chituc *et al.* 2016: 22).

Using the same measurement scale from Experiment 1, participants report the extent to which they agree with the following three statements: (i) at 11:45 AM, Adams ought to keep his promise, (ii) Adams can keep his promise, and (iii) Adams is to blame for not keeping his promise. Chituc *et al.* found a modest correlation between *ought* and *blame* judgments, but did not find such a correlation between *ought* and *can* judgments. Moreover, after conducting a further analysis of the study that excluded 29 participants (out of 195 participants) who apparently misread the vignette by providing *can* ratings above the midpoint, Chituc *et al.* found that *ought* was strongly correlated with *blame*, but not with *can*. They find these results to cast doubt on the blame validation hypothesis with respect to Experiment 1; if participants who are inclined to accept OIC are subject to the process of blame validation, then, contrary to the results of Experiment 2, we should find a strong correlation between *ought* and *can*.

But there are two reasons to doubt that the results of Experiment 2 undercut the blame validation hypothesis. First, Adams is presumably *not* blameworthy for failing to keep his promise since his car breaks down *unexpectedly*. So all *blame* ratings that are above the midpoint suggest that such participants failed to comprehend the vignette.⁶ Consequently, any correlations with *blame* ratings that are above the midpoint are suspect. This is because the blame validation hypothesis, as I understand it, is a hypothesis about the responses of participants who are competent enough to understand basic facts about the vignette, even if such participants are not competent enough to avoid forming muddled intuitions.

Second, the correlations between the *blame* ratings and *ought* ratings that are significantly below the midpoint imply that such participants are not subject to the process of blame validation precisely because they do not find Adams blameworthy. But this in no way shows that participants in Experiment 1 who *do* find Adams blameworthy in *High Blame* are *not* subject to the process of blame validation. So any correlations with *blame* ratings that are significantly below the midpoint appear to have no evidential bearing upon the blame validation hypothesis.

Putting aside the blame validation hypothesis, Chituc *et al.* also consider the absence of a correlation between *ought* and *can* to count against OIC. The evidential bearing of this absence for OIC is also challenged by the *temporal dimension* hypothesis. Recall that Hannon's study supports the hypothesis that Chituc *et al.*'s *High Blame* question did not highlight to a sufficient extent the distinction between the pre-inability and post-inability 'oughts'. The same issue arises for the *ought* claim in Experiment 2, especially since this claim lacks the word "still" which plays a role in highlighting to some extent the distinction between the pre-inability and post-inability 'oughts'. So

⁶ Since the mean rating for *can* ($M = -33.01$) is lower than the mean rating for *blame* ($M = -15.48$), and since 29 participants provided ratings for *can* that are above the midpoint, we have some reason to think that more than 29 participants provided ratings for *blame* that are above the midpoint.

Experiment 2 does not appear to strongly count against OIC or against the blame validation hypothesis.

6.

An additional method for evaluating the results of these empirical studies involves identifying an alternative principle to OIC that commonsense morality may endorse. Buckwalter and Turri (2015) suggest that commonsense morality may instead accept a “blame implies can” (BIC) principle in light of the results of one their studies (Experiment 8). In this study participants read the ‘car accident’ version of the vignette from Buckwalter and Turri’s Experiment 1. Participants in one group select one of the four answers from Experiment 1 in order to test OIC. Participants in the other group select one of the following four options in order to test BIC:

1. Walter is blameworthy for not picking up Brown at the airport, and Walter is physically able to pick him up.
2. Walter is blameworthy for not picking up Brown at the airport, but Walter is not physically able to pick him up.
3. Walter is not blameworthy for not picking up Brown at the airport, but Walter is physically able to pick him up.
4. Walter is not blameworthy for not picking up Brown at the airport, and Walter is not physically able to pick him up.

While 85% of participants in the first group selected the “Ought but Unable” response, 97% of participants selected the “Not blameworthy and Unable” response. In light of the stronger correlation between “blame” and “can”, Buckwalter and Turri (2015: 15) suggest that commonsense morality endorses BIC rather than OIC.

We need a precise formulation of BIC in order to compare it to OIC. Since this paper’s formulation of OIC indexes both an agent’s obligation and her ability to the same time, we might attempt a similar formulation for BIC:

Necessarily, if at time t an agent A is blameworthy for ϕ -ing, then at t A can ϕ .

This formulation does not appear to be what Buckwalter and Turri have in mind. It implies that an assassin who is presently blameworthy for killing the king *can* presently kill the king. In order to find a better formulation, we need to look for one that gains evidential support from the results of Experiment 8. Under the assumption that “Not blameworthy and Unable” is the correct response, BIC will not apply to Experiment 8; since Walter is *not* blameworthy for breaking his promise, we cannot deduce anything about Walter’s abilities on the basis of BIC. The contrasting vignette to Experiment 8 appears to be Henne *et al.*’s *High Blame* in which Adams is (let’s stipulate for the sake of argument) blameworthy for not picking up Brown at noon as promised. This seems to be the sort of vignette to which BIC applies. Since Adams is blameworthy for not picking up Brown at noon, we must be able to deduce an ability that Adams possesses on the basis of BIC. So perhaps the formulation that we are after goes like this:

Necessarily, if at time t an agent A is blameworthy for not ϕ -ing, then at t A can ϕ .

But this principle implies that a lifeguard who is presently blameworthy for not saving a child who drowned an hour ago can (at the present moment) still save that child. It also implies that a retired Governor who is blameworthy for not signing a bill twenty years ago can still sign that bill. If these formulations are not what Buckwalter and Turri have in mind, then it's unclear what is being compared to OIC. If they do have in mind one of these formulations, then OIC remains the more plausible principle.

Henne *et al.* (2016) also suggest a contender to OIC. When considering a case in which an agent, Brit, cannot meet Dan for a beer as promised due to an automobile accident, Henne *et al.* suggest that most people would retract the claim that Brit ought to meet Dan, not simply because she can't meet Dan, but because Brit is not blameworthy for being unable to meet Dan. So Henne *et al.* seem to support the following principle:

“Blame” Implies “Ought” (BIO): Necessarily, if at time t an agent A is blameworthy for losing an ability to fulfill an obligation to ϕ , then at t A still has an obligation to ϕ .

This principle implies that a lifeguard who is presently blameworthy for losing an ability to save a child who drowned an hour ago has (at the present moment) an obligation to save that child. It also implies that a retired Governor who is blameworthy for losing an ability to sign a bill twenty years ago still has an obligation to sign that bill. Given these implications, OIC also appears to be more plausible than this contender.⁷

7.

This paper has defended the position that the current empirical data on OIC does not show that OIC is not intuitive, or that OIC is false, or that OIC is less plausible than BIC or BIO. To conclude, this paper will consider a final anti-OIC argument by Henne *et al.* that is primarily based upon the results of Chituc *et al.* The argument goes like this:

1. If OIC is true, it is a conceptual truth.
2. OIC is not a conceptual truth.
Therefore,
3. OIC is false.

The evidence for premise (2) is based upon the studies by Buckwalter and Turri (2015) and Chituc *et al.* (2016). But if the *temporal dimension* hypothesis is correct, or if the blame validation hypothesis is correct (or both since they are not inconsistent with one another), then the evidence for (2) seems to be undercut. Moreover, defenders of OIC such as Michael Zimmerman (1996) and Peter Vranas (2007) who accept (1) (or something close enough) are unlikely to be surprised by the results of these studies since they admit that conceptual or analytic truths like OIC can be controversial and opaque:

“Indeed, [OIC] will be analytic; but its being so is no reason to expect it to be

⁷ Henne *et al.* (2016: 288) say that most people's retraction of the claim that Brit ought to meet Dan suggests that 'ought' implies 'blame', and that this principle accords with Chituc *et al.*'s results. It is likely that they meant the converse since Chituc *et al.*'s results seem to support the converse principle, and since it is possible for someone to be obligated to do something without being blameworthy for anything.

uncontroversial...” (Zimmerman 1996: 77).

“Conceptual entailment need not be transparent...So a version of the ought-implies-can principle formulated in terms of conceptual entailment need not be indisputable or insignificant if it is true” (Vranas 2007: 170).

The empirical results of Kurthy *et al.* and Hannon confirm the view that analytic or conceptual truths can be opaque, e.g. when non-philosophers unwittingly select responses that rationally commit them to the denial of OIC.

When the responses in these experimental studies do not highlight to a sufficient extent the distinction between the pre-inability and post-inability ‘oughts’, participants are motivated to select the answer that best captures the temporal dimension of the story in order to display their comprehension of that story. Since the response that best captures the temporal dimension of the story in certain experimental studies is inconsistent with OIC, participants form muddled intuitions by focusing upon the pre-inability ‘ought’ rather than the post-inability ‘ought’. Moreover, if the blame validation hypothesis is true with respect to some participants, they are motivated to seek out a justification for attributing blame to the agent, and thus select the answer that accompanies the muddled intuition that the agent ought to have done something that they could have done at the time, even though the selected answer concerns a post-inability ‘ought’. Given the viability of these hypotheses, the debate over OIC is far from settled.⁸

References

- Alicke, M. 2000. Culpable control and the psychology of blame. *Psychological Bulletin* 126 (4): 556–574.
- Buckwalter, W. and J. Turri. 2015. Inability and obligation in moral judgment. *PLOS ONE* 10(8): 1–20.
- 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.
- Graham, P. A. 2011. ‘Ought’ and ability. *Philosophical Review* 120 (3): 337–382.
- Greene, J., Sommerville, R., Nystrom, L., Darley, J., and Cohen, J. 2001. An fMRI investigation of emotional engagement in moral judgment. *Science* 293 (5537): 2105–2108.
- Hannon, M. Forthcoming. Intuitions, reflective judgments, and experimental philosophy. *Synthese*. DOI 10.1007/s11229-017-1412-1
- Henne, P., Chituc, V., De Brigard, F., and W. Sinnott-Armstrong. 2016. An empirical refutation of “ought” implies “can.” *Analysis* 76 (3): 283–290.
- Kurthy, M. and H. Lawford-Smith. 2015. A brief note on the ambiguity of ‘ought’. Reply to Moti Mizrahi’s ‘Ought, can and presupposition: an experimental study’. *Method* 4 (6): 244–249.
- Kurthy, M., Lawford-Smith, H., Sousa, P. 2017. Does ought imply can? *PLOS ONE* 12 (4): 1–24.
- Mizrahi, M. 2015a. Ought, can, and presupposition: an experimental study. *Method* 4 (6):232–243.

⁸ Thanks to Neil Feit, Paul Henne, David Hershenov, Stephen Kershner, Alex King, David Limbaugh, Jake Monaghan, Travis Timmerman, the audience at the 2017 ‘Blameless Buffalo?’ conference, and two referees for this journal for helpful feedback on previous drafts of this paper.

- Mizrahi, M. 2015b. Ought, can, and presupposition: A reply to Kurthy and Lawford-Smith. *Methodes* 4 (6): 250–256.
- Nichols, S. 2002. Norms with Feeling. *Cognition* 84 (2): 221–236.
- Nichols, S. and J. Knobe. 2007. Moral responsibility and determinism: the cognitive science of folk intuitions. *Noûs* 41(4): 663–685.
- Schroer, J. and R. Schroer. 2013. Two potential problems with philosophical intuitions: Muddled intuitions and biased intuitions *Philosophia* 41(4): 1263–1281.
- Vranas, P. 2007. I ought, therefore I can. *Philosophical Studies* 136(b): 167–216.
- Zimmerman, M. 1996. *The Concept of Moral Obligation*. Cambridge: Cambridge University Press.