

Ability, Responsibility, and Global Justice *

Wesley Buckwalter

wesleybuckwalter@gmail.com

Abstract: Many have argued we have a moral obligation to assist others in need, but given the scope of global suffering, how far does this obligation extend? According to one traditional philosophical view, the obligation to help others is limited by our ability to help them, or by the principle that “ought implies can”. This view is primarily defended on the grounds that it is a core principle of commonsense moral psychology. This paper reviews findings from experimental philosophy in cognitive science demonstrating that “ought implies can” is rejected by moral psychology. Researchers find that moral obligations are ascribed to agents who cannot fulfill them, suggesting that moral requirements do sometimes extend beyond what we are able to do. This research furthers our understanding of moral obligation, identifies an important need for further cross-cultural work in moral psychology, and demonstrates a way in which scientific experimentation can be applied to improve upon the conceptual analysis of important philosophical concepts in normative ethics.

Keywords: ethics; assistance; demandingness; moral judgment; ability; obligation; “ought implies can”

* This is the penultimate version of a paper forthcoming in *Journal of Indian Council of Philosophical Research*. Please cite the published version.

Introduction

It is estimated that over seven hundred million people live in extreme poverty worldwide (Cruz, Foster, Quillin, and Schellekens 2015). Over sixty-five million have been forcibly displaced from their homes due to violence, persecution, or human rights violations. Twenty million of these displaced persons are refugees and half of those refugees are children (UNHCR 2016). Each day over twenty thousand children under the age of five die largely because of hunger and preventable disease (UNICEF 2014). Natural disasters such as earthquakes, heat waves, and landslides affect nearly one hundred million people and did sixty-six billion dollars worth of damage around the world in one year alone (UNISDR 2015).

Many have argued that we have moral responsibilities to assist those in need by donating to global assistance organizations, opening our borders to refugees, opposing human rights violations, and supporting disaster relief programs (Sen 1999; Singer 1972; Carens 2013; Kamm 2008). However it is a matter of debate how far obligations to help others in need extend (Shue 2004; O'Neill 2004, 2016; Miller 2007; Arneson 2004; Hurley 2003). Clearly the average human being could not resolve global poverty or prevent even a small fraction of widespread human suffering. It is unlikely that any one country or group could single handedly relocate all refugees or displaced persons. And often, practical challenges due to international borders, geographic features, and decimated infrastructures undermine humanitarian assistance efforts or make delivering aid to remote areas of the world effectively impossible. Given these realities, what is the limit to our moral responsibility to ease human suffering and how much does one need to do before this obligation counts as satisfied?

According to one answer to this question, having a moral obligation to help others in need requires having the ability to help others in need. In other words, the morally correct thing to do in any one situation never exceeds what one is able to do. On this theory, for example, the maximum amount of money a person is morally required to donate or the number of refugees a country is morally required to accept is fully determined by the amount of money they are able to give or the number of persons they are able to accept. Similarly, on this view, one cannot be morally obligated to help others in remote areas of the world if effectively delivering aid to those areas is impossible.

According to another answer to this question, the moral obligation to help those in need is not always constrained by our ability to help them. In other words, the morally correct thing to do can, in some cases, exceed what one is able to do. On this view, for example, the resources that a person or a country has can impact the amount of money they end up donating or the number of refugees they end up accepting, but these financial facts do not fully answer the normative question concerning what morality requires. Similarly, on this view, one can have a moral obligation to help others in need in remote areas of the world, despite how difficult it may be to deliver aid to that area effectively.

The view that obligations to promote global justice, such as helping others in need depends on our abilities is motivated by the longstanding ethical principle that “ought implies can”. Support for this principle weaves its way through the history of western philosophy and is accepted by most contemporary anglophone philosophers (Kant 1998; Littlejohn 2012; Moore 1922; Hare 1965; Van Fraassen 1973; Dahl 1974; Vranas 2007; Copp 2008; Cicero and Edmonds 1856). Although the principle is glossed in many forms, the basic idea is that inability rules out obligation. According to the principle, an agent

cannot be morally required to act if they are unable to perform that action. Applying this principle to global aid obligations, Onora O'Neill writes "institutions and individuals can have obligations *if but only if* they have adequate capabilities to fulfill or discharge those obligations," and subsequently, that "individuals cannot be obliged to resolve the problems of world hunger, or to grow wings and fly; institutions cannot have obligations to perform tasks for which they lack capabilities" (2004: 251). If individuals or institutions are unable to perform an action, such as assist others in need, then morality does not require that they assist them.

One of the main arguments offered in support of "ought implies can" is that the principle best explains core features of moral psychology. For example, some theorists claim that the principle is simply "a clear and evident truth" of moral cognition that enjoys "widespread assent" among theorists and ordinary people (Stocker 1971: 303; Huemer, as quoted by Ryan 2003: 57). Other theorists claim that the principle best explains patterns of moral discourse observed during the give and take of ordinary moral exchanges. For example, some theorists claim that one "*cannot say* of anyone that he ought to do a certain thing, if it is a thing which it is physically impossible for him to do," (Moore 1922: 317, emphasis added), or that the "lack of capability always counts against *an ascription* of obligations" (O'Neill 2004: 251, emphasis added). However these arguments are not without critics. For example, some theorists have denied these claims, arguing that there are counterexamples to the principle in which obligation is attributed to agents without ability (Sinnott-Armstrong 1984; Ryan 2003; Graham 2011). These arguments question the main support offered for the "ought implies can" principle.

Given this disagreement, and given that the main argument for the principle in question is largely psychological in nature, some philosophers have turned to the tools of experimental philosophy in cognitive science to investigate support for “ought implies can”. Rather than rely only on introspection about what principles might appear obvious, or on casual observation about social practices, these researchers also conduct controlled experimental studies in an attempt to study moral judgment with greater rigor and precision than speculative methods alone can provide (for reviews, see Knobe et al. 2012; Sarkissian 2016; Buckwalter and Turri Forthcoming). To date, several teams of researchers have applied experimental methods to study patterns of obligation ascriptions and to see whether those patterns actually provide evidence for the “ought implies can” principle (Buckwalter and Turri 2015; Mizrahi 2015b, 2015a; Chituc, Henne, Sinnott-Armstrong, and De Brigard 2016; Turri 2017b, 2017a). When they did this, the researchers found that moral obligations are frequently ascribed to agents who are unable to fulfill them. In other words, the evidence each team of researchers uncovered overwhelmingly indicates that commonsense moral cognition, at least among anglophones, utterly rejects ought-implies-can.

The remainder of this paper reviews the discoveries that researchers have made concerning the effects of ability and inability on moral judgment and discusses how these findings advance our philosophical understanding of moral obligations. The findings have significant implications for theories of aid obligations in discussions of global justice specifically and for the traditional role that ability has played in ethical theorizing generally. The findings also serve as a case in point for one way in which experimental

techniques can be used as a tool to aid and improve upon the conceptual analysis of important philosophical concepts in normative ethics.

Experimental Evidence

This section reviews experimental findings from several teams of independent researchers questioning the entailment between moral obligation and ability in moral judgment. Using a variety of different experimental paradigms, narrative cover stories, probing methods, and statistical techniques, each team of researchers provides consistent and convergent evidence that “ought implies can” is a false principle of commonsense moral psychology.

Some of the earliest experimental evidence against the “ought implies can” principle in moral psychology investigated moral ascriptions to agents in straightforward cases of moral obligation and physical disability (Buckwalter and Turri 2014, 2015). In one study, for example, researchers presented participants with the following stimuli involving the obligation to help others in need (Buckwalter and Turri 2015: Experiment 7). The basic story began like this:

Michael is relaxing in the park when he sees a small girl fall into a nearby pond. She is drowning and definitely will die unless someone quickly pulls her out. This part of the park is secluded and Michael is the only person around.

Researchers then manipulated whether the protagonist had an ability to act in this case. Roughly half of the participants saw the case continue in a way that made it clear Michael had the ability to act to save the victim:

And Michael is a normal adult male and can swim fast enough to save the girl. As a result, Michael is physically able to save the girl.

The other half of the participants saw the case continue in a way that made it clear that Michael lacked the ability save the victim:

But Michael is stricken with a sudden paralysis in his legs and cannot swim to save the girl. As a result, Michael is not physically able to save the girl.

Researchers then asked participants a series of questions designed to gauge what abilities Michael had, what his moral obligations were, and whether Michael deserves blame for the outcome in the story. For example, in order to check that participants understood the cases, researchers asked participants whether they thought that “Michael is literally unable to save the girl”. The results indicated that participants did understand the cases. Nearly all participants answered “no” when Michael was described as having the ability and nearly all answered “yes” when Michael was described as suffering from paralysis.

Next, researchers asked whether Michael had a moral obligation in light of his ability or inability to act. Prior research indicated that biases and other pragmatic pressures can occur when applying concepts like obligation, ability, and blame in isolation (see Alicke 2008 on "blame validation"; Turri and Blouw 2015 on "excuse validation"). For this reason, researchers took several steps to minimize potential biases in measuring obligation judgments. To do this, researchers designed a novel conceptual combination task, which forced participants to consider the concepts of ability and moral obligation simultaneously and to make a single judgment about them. To further reduce the possibility of biased responses, researchers chose simple connectives such as “and” or “but” to express conjunctions between concepts instead of using other, more complicated

logical connectives. Guided by this approach, researchers asked participants to select the one response option that best describes Michael's situation concerning his obligations and abilities:

1. Michael is morally obligated to save the girl, and Michael is physically able to do so.
2. Michael is morally obligated to save the girl, but Michael is not physically able to do so.
3. Michael is not morally obligated to save the girl, but Michael is physically able to do so.
4. Michael is not morally obligated to save the girl, and Michael is not physically able to do so.

Researchers found that 83% of participants selected the “morally obligated and able” response option when Michael was described as being a normal adult male, while 78% of participants selected the “morally obligated but unable” response option in the case of paralysis. The overwhelming majority of participants indicated that Michael had the moral obligation to save the girl, while at the same time directly acknowledging that he either did or did not have the physical ability to do so. In other words, facts about the protagonist's ability to act had no measurable impact in this circumstance on whether that moral obligation was judged to be present. Moral obligation was ascribed at high rates despite ability or inability.

Lastly, researchers asked participants to rate their level of agreement with the following statement regarding blame for the outcome: “Michael deserves to be blamed for the fact that the girl drowns.” Researchers found that facts about ability strongly

impacted blame judgments for this result between the two cases. Participants denied that Michael deserves to be blamed for the outcome when he is unable to act. Participants were ambivalent, however, when Michael did have ability and failed to save the victim. In this case, participants blamed Michael for the fact the girl drown at approximately chance rates. In other words, while ability did not impact obligation judgments in this case, it did have a large impact on blame attributions (see also Chituc, Henne, Sinnott-Armstrong, and De Brigard 2016; Turri and Blouw 2015; Turri 2017a: for related results and discussion).

Researchers replicated these same basic results across a wide range of cover stories and narrative contexts featuring many different types of obligation and different kinds of disability. In other experiments, for instance, to avoid a possible objection about temporal ambiguity in the stimuli, researchers manipulated the scope and temporal duration of the disabilities that rendered agents unable to act (Buckwalter and Turri 2015: Experiments 5-6). For example, researchers showed that moral obligation ascriptions persist even when Michael had a lifelong disability (e.g. he was “paralyzed since birth”) or when the circumstances would have generally prevented all agents from acting (e.g. because “no human could” have saved the victim). One plausible alternative explanation of the finding above is that when participants ascribe moral obligation to Michael, they only mean to do so up until the point in the story when he becomes unable to act, and not after, which would be consistent with “ought implies can”. But the present results strongly rule out the possibility that any temporal ambiguity occurred. Participants ascribe moral obligations to characters without the abilities to fulfill them at high rates even though there was no point in the stimuli when they had the ability to fulfill them. At

the same time, participants also continued to deny blame to agents for not fulfilling their obligations when inability prevented from doing so.

Researchers replicated and expanded upon these effects in several ways. For example, they also investigated different kinds of obligations, ranging from promises to general duties to prevent harm, and severity of the consequences for failure, ranging from mundane inconvenience to the tragic loss of life, to see what affect these things might have on moral obligation ascriptions. Again and again, researchers found that moral obligations were continually ascribed to agents despite their inability to fulfill them in each of these circumstances. Researchers also repeated the same experiments using different moral terminology to probe moral judgments. For example, instead of asking about the “moral obligation” a protagonist had, researchers also collected judgments about their “moral duties”, what the protagonists “morally ought” to do, or “had a moral responsibility” to do (see also Turri 2017a). Probing terminology made no detectable difference to the patterns of responses observed by the researchers. Moral obligation, moral duties, and moral oughts were each consistently judged to persist beyond ability in the cases tested.

Around the same time, other groups of researchers found convergent evidence for the same judgments that moral requirements persist despite inability. In one series of papers, for example, researchers found that obligations from prior commitments are not, or at least, not always nullified just because unfortunate circumstances prevent one from fulfilling them (Mizrahi 2015a, 2015b; though also see response by Kurthy and Lawford-Smith 2015). In this study, participants read a story about Smith, a professor who commits to meeting with her student Sid during office hours later that day. However the

professor becomes unable to keep his office appointment with the student, either because he “gets locked in his classroom before he is able to make it to his office” or because he simply “forgets about having office hours that day”. In both cases, participants disagreed with the statement “Professor Smith can keep his office appointment with Sid.” At the same time, however, participants agreed with the statement “Professor Smith ought to keep his office appointment with Sid.” In other words, judgments about what one “ought” to do outstripped judgments about what they “can” do, which is the complete opposite pattern of judgments predicted by the claim that “ought implies can” is an intuitive principle of ordinary moral judgment. These findings also replicate and extend other results concerning the “location” of an inability to act (Buckwalter and Turri 2015: Experiment 4). Specifically, the present results show that patterns of judgments systematically violate the “ought implies can” principle when inabilities result from factors both “outside” an agent, due to opportunity or physical location, like being locked in a classroom, and “inside” the agent, due to disablement of an agent’s personal force or due to simple forgetfulness.

Another piece of converging evidence against this principle comes from research on obligation and blame judgments for failed promises (Chituc, Henne, Sinnott-Armstrong, and De Brigard 2016; Henne, Chituc, De Brigard, and Sinnott-Armstrong 2016). Across several narrative contexts, these researchers show that participants judge that agents ought to honor promises even after they are no longer able to keep them. In one experiment, for example, researchers present participants with stories that involve a promise to see a movie with a friend (Chituc, Henne, Sinnott-Armstrong, and De Brigard 2016: Experiment 3):

Brown is excited about a new movie that is playing at the cinema across town. He hasn't had a chance to see it, but the latest showing is at 6 o'clock that evening. Brown's friend, Adams, asks Brown to see the movie with him, and Brown promises to meet Adams there. It takes Brown fifteen minutes to drive to the cinema, park, purchase a ticket, and enter the movie. It would take 30 min if Brown decided to ride his bike. The cinema has a strict policy of not admitting anyone after the movie starts, and the movie always starts right on time.

The story then continued in one of two ways that made it clear that Brown cannot make it to the cinema in time, thereby rendering him unable to fulfill his promise to see the movie with Adams. In one version of the story, the "low fault" condition, Brown was unable to attend the movie through no fault of his own:

At 5:30, Brown thinks about riding his bike, but decides it is too cold. Instead, he leaves at 5:45, but his car breaks down five minutes later. He can't fix it himself in time to make it to the cinema, and it is too late to make it by bike. Because his car is not working at the time, Brown can't meet his friend Adams at the movie by 6.

In a separate version of the story, the "high fault" condition, Brown was unable to attend the movie directly through a fault of his own:

As Brown gets ready to leave at 5:45, he decides he really doesn't want to see the movie after all. He passes the time for five minutes, so that he will be unable to make it to the cinema on time. Because Brown decides to wait, Brown can't make it to the movie by 6.

After seeing one of these cases, participants rated whether “At 5:50, Brown can make it to the theater by 6”, whether “Brown ought to make it to the theater by 6”, and whether “Brown is to blame for not making it to the theater by 6”.

Researchers found that participants strongly disagreed with the “can” attribution in both cases. They found that participants tended to agree with the “ought” attribution in the high fault condition and that they tended to disagree in the low fault condition. And similarly, participants tended to agree with the “blame” attribution in the high blame condition and that they tended to disagree in the low blame condition. When testing the relationship between these judgments, researchers found that while there was a significant correlation between “ought” and “blame” judgments, there was no correlation between “ought” and “can” judgments. If the “ought implies can” principle is supported in moral psychology, we would predict a strong correlation between these judgments. But no correlation was found. On the basis of these findings, then, the researchers conclude that there is no analytic or conceptual entailment in moral psychology between the concepts expressed by “ought” and “can”.

Another piece of convergent evidence against “ought implies can” comes from research on free will judgments in metaphysics and the question of whether determinism is compatible or incompatible with moral responsibility (Turri 2017b). In these experiments, researchers showed that moral responsibility is ascribed to protagonists in situations where it is literally impossible for them to act otherwise from how they did. In one experiment, for example, researchers presented participants with cases that manipulated whether a certain outcome was possible or impossible (Turri 2017b: Experiment 2). The differences between cases are indicated in the brackets below:

A woman is evaluating her employee's performance. The employee performed excellently. Given the current condition of the woman's brain, it is [impossible/possible] for her to give the employee a good evaluation. As a matter of brain chemistry, it is [literally impossible/possible] that she can give the employee a good evaluation. She does not give the employee a good evaluation.

Participants were asked, among other things, whether they agreed or disagreed that "The woman could have given the employee a good evaluation", whether she "chose not to give the employee a good evaluation", and whether she "had a moral responsibility to give the employee a good evaluation."

Researchers found extremely large differences in "could" judgments between possibility and impossibility conditions, indicating that they understood and accepted the conditions stipulated in the stories (for discussion of comprehension of free will scenarios, see Rose, Buckwalter, and Nichols 2017). Researchers also detected large differences in "choice" judgments between conditions. Participants strongly agreed that the protagonist chose not to give a good evaluation in the possibility condition and disagreed that a choice was made in the impossibility condition. Lastly, researchers found that moral responsibility was strongly ascribed in both possibility and impossibility conditions. The facts that the protagonists did not make a choice or that it was literally impossible for the protagonist to give a good evaluation made no difference to judgments that the protagonist was morally responsible. Researchers replicated this same basic finding across multiple cover stories and probing methods.

These findings provide additional evidence questioning the support for "ought implies can". According to this principle, one cannot have a moral responsibility to do

something when its performance is impossible. Thus, if “ought implies can” is a principle of moral psychology, we should predict both large differences in responsibility ascriptions between possibility and impossibility conditions, and also that moral responsibility is denied at greater than chance rates in the impossibility condition. But in reality, neither of these things happened. Moral responsibilities were strongly ascribed to the protagonist independently of the possibility to do otherwise. Moral responsibility for behavior was thought to persist even when the action required was literally impossible.

Lastly, a final piece of evidence against “ought implies can” comes from observations involving excuses and blame. When we fail to act in certain ways, we often attempt to offer excuses to explain our behavior. Often, the excuses we offer appeal to facts that pertain to inability, such as external forces, manipulation, situational factors, or outcomes beyond our control. This behavior might be explained in one of two ways. One explanation, consistent with “ought implies can”, is that we offer these kinds of excuses in order to signal that we were not morally responsible or obligated to act otherwise. After all, if “ought implies can” is true and obligation entails ability, then what is there one needs to be excused from doing in such cases in the first place? A rival explanation, inconsistent with “ought implies can”, is that we offer excuses involving inability to exculpate the blame that typically accompanies unfulfilled obligations. This explanation challenges “ought implies can” because it posits obligation despite inability.

In line with the latter explanation, researchers have found that judgments about excuses constitute additional evidence against “ought implies can” in moral psychology (Turri 2017a). In one study for example, researchers presented participants with cases very similar to the ones above involving the employer who was unable to give the

deserving employee a positive evaluation (Turri 2017a: Experiment 1). Researchers then asked participants to choose which of two response options best describes that case. The option consistent with “ought implies can” was:

The inability prevents her from having a moral obligation to give the employee a positive evaluation. There is nothing to excuse because she did not have an obligation.

The option inconsistent with “ought implies can” was:

The inability prevents her from being blameworthy for not giving the employee a positive evaluation. The inability is a legitimate excuse for not fulfilling an obligation.

Researchers found that the overwhelming majority of participants chose the second option. The participants indicated that the protagonist had an obligation to give the employee a positive evaluation, but that her inability to do so should excuse her from being blamed for failing to fulfill that obligation.

In subsequent studies, researchers replicated and expanded upon this effect in several ways. For example, researchers replicated the same pattern of responses reported above using different probing methods, such as a “check all that applies” task. Using causal search algorithms and structural equation modeling, research also showed that blame judgments in similar kinds of cases were fully mediated by judgments about excuses involving the inability to act, independently of judgments about the presence of moral obligation (Turri 2017a: Experiment 3). This finding suggests that inability can constitute a legitimate excuse that renders agents blameless for failing to do something that they were nonetheless morally obligated or responsible for doing. This research

further our understanding of excuses and their connection to blameworthiness. For example, it suggests that blamelessness for failed obligations is sometimes explained by having a legitimate excuse and inability can be a legitimate excuse. It also constitutes further evidence that obligation can sometimes persist despite inability in moral psychology, even if blameworthiness for failing to fulfill obligations in that circumstance may not.

Subsequent studies on the connection between moral judgment and ability also demonstrate contexts in which ability is related to obligation ascriptions. In one study, for example, a team of researchers presented participants with cover stories modeled after prior materials developed by Buckwalter & Turri 2015 involving a protagonist who is unable to fulfill a promise due to a car accident (Kurthy, Lawford-Smith, and Sousa 2017: Experiment 1). The researchers replicated the previous finding that a large majority of participants indicate that the protagonist was “obligated but unable” to act. However the researchers also demonstrate that when probed in a different manner, a large majority of participants indicate that the protagonist was “not obligated because they were unable” to act. The researchers replicated this effect across a series of studies. These findings further advance our understanding of moral judgments in two ways. First, consistent with past work they continue to suggest that conceptual entailment between obligation and ability is false in commonsense judgment. Second, they suggest that although entailment is false there could nonetheless be important inferential connections between moral judgments and ability ascriptions and that discovering the nature of these connections could improve our understanding of commonsense “ought” and “can” judgments.

Conclusions for Ethical Theory

The “ought implies can” principle is the traditional philosophical principle that moral obligation is incompatible with inability. The view places limits on the things that morality requires of us in different situations. Applied to discussions of global justice, for example, the view says that individuals and institutions cannot be morally obligated to help others in need, accept refugees, or donate to humanitarian assistance programs when they are unable to do so. Support for the principle in ethics is often given on the grounds that “ought implies can” is a principle of commonsense moral cognition that reflects patterns of obligation ascription. When researchers put the principle to the test, however, they found that the patterns of ascriptions do not support “ought implies can”. Across several narrative contexts, experimental designs, probing methods, and statistical techniques, researchers found that moral obligations, duties, and responsibilities are attributed to agents in cases where ability to fulfill them is absent. For example, participants judged that an innocent bystander has the moral obligation to save someone in need, regardless of whether the bystander has the ability to do so. These results demonstrate that moral judgments are not always limited by ability and thus that it is false that “ought implies can” is a principle of commonsense moral psychology.

Though the results falsify the claim that “ought implies can” is a core principle of commonsense moral cognition, they also demonstrate some other ways in which ability is highly relevant to moral judgment. For example, though perceptions of inability did not deter participants from ascribing moral obligation at very high rates in several narrative contexts, inability did have a large impact on ascriptions of blame for unfilled obligations. Across several experiments, researchers found that participants were much

less likely to blame agents for failure to fulfill an obligation when the reason for failure was explained by a physical disability due to paralysis or a car crash. Subsequent experimental research also revealed a psychological pathway in which these judgments are made. The research suggests that agents are often found to be blameless when they are unable to help others, not because they did not have an obligation to do so in that circumstance, but rather because inability can be a legitimate excuse for failing to fulfill moral obligations that persist after one becomes unable to act.

These results contribute to our understanding of moral obligations in normative ethics in several ways. The results undermine the central theoretical motivation for accepting “ought implies can”. Thus, the argument that the principle should be accepted because it is a principle of commonsense moral cognition fails. The results also provide some reason to think the principle itself is false. Several theorists have argued that theorizing in ethics should acknowledge and reflect core principles of moral psychology. For example, Michael Brownstein writes that, “a good theory of moral responsibility should be relatively consistent with—or make sense of—common folk attitudes toward responsibility, since the concept of responsibility is itself a deeply social one” (2016: 3). This sentiment nicely captures the motivation for one of the main arguments for “ought implies can” in the first place, which relied on patterns of judgments in ordinary thought and talk to support the principle. Subsequent research showed that theorists relying on this type of evidence were wrong about what those patterns were. There is strong evidence for the systematic ascription of moral obligation to agents who are unable to fulfill them.

The findings also have implications for discussions of global aid obligations. Given that obligation can persist despite inability in moral judgment, it is unlikely that facts about the finances of persons or countries will fully answer the question of what the morally correct thing to do is in cases of extreme need. The moral obligation to help others, may, in some circumstances, extend beyond the ability to help them. At the same time, however, inability may provide a legitimate excuse for why some moral obligations to help others remain unfulfilled and why individuals may be blameless for failing to do so. Further research is needed to study these judgments in contexts of development and assistance specifically. If the “ought implies can” principle is false, however, it removes a major theoretical challenge to the view that robust moral obligations to promote global justice exist despite the level or extent of extreme need.

The results also demonstrate how experimental investigation of moral judgments can improve upon the conceptual analysis of important philosophical concepts in normative ethics. For example, some philosophers have thought that having a moral obligation or responsibility is essentially the same thing as being blameworthy for failing to do it (e.g. Mill 1979). But the present results point to a conceptual distinction between moral responsibility and blame. The experiments demonstrate large and systematic differences in responsibility and blame ascriptions to agents in circumstances in which they were able and unable to act. In some experiments, responsibility judgments correlated with blame judgments. In other experiments, responsibility was ascribed in cases where blame was not. This suggests that there is an important psychological difference between the processes used to identify the presence of a moral obligation an agent has, on the one hand, and those used to evaluate agents when they fail to fulfill those obligations, on the

other. The fact that there are some circumstances in which agents have unfulfilled moral obligations but are not considered blameworthy also supports and expands upon the often-overlooked moral category of “blameless transgression” (Turri 2013; Turri and Blouw 2015; see also Cushman, Durwin, and Lively 2012). Future research in philosophy and cognitive science might build on this discovery to further explore the contours of this subtle conceptual distinction and moral category.

Lastly, this approach identifies an important need for cross-cultural research experimentally investigating ascription practices among native speakers of different languages. The main argument in favor of “ought implies can” appeals to moral psychology, but few theorists have considered whether such support is universal or culturally local. The experiments reported above were conducted in the United States among English speakers. But it is an open question whether the same patterns of moral judgments exist in other cultures (see Doris and Plakias 2008; Graham, Meindl, Beall, Johnson, and Zhang 2016). The current evidence for this claim is mixed. For example, researchers have found that both native English and German speakers are sometimes willing to ascribe obligations to agents for promises they are unable to keep, further questioning “ought implies can” (Willemsen 2016: Experiments 1-2). In the same studies, however, researchers also found that blame judgments varied across cultures. English speakers were more willing to blame agents for failing to fulfill obligations than German speakers were. Thus it would be valuable to test whether patterns of moral judgment support or continue to violate “ought implies can” across a wider array of cultures and languages, as well as whether these judgments differ from the evaluations of agents for unfulfilled obligations through blame or punishment.

Acknowledgements – For helpful feedback I thank Carolyn Buckwalter, Stephen Stich, and John Turri. This research was supported by a Banting Fellowship awarded through the Social Sciences and Humanities Research Council of Canada.

References

- Alicke, Mark (2008), 'Blaming Badly', *Journal of Cognition and Culture*, 8 (1), 179-86.
- Arneson, Richard J. (2004), 'Moral Limits on the Demands of Beneficence?', in Deen K. Chatterjee (ed.), *The Ethics of Assistance: Morality, Affluence, and the Distant Needy* (Cambridge: Cambridge University Press), 33-58.
- Brownstein, Michael (2016), 'Attributionism and Moral Responsibility for Implicit Bias', *Review of Philosophy and Psychology*, 7 (4), 765–86.
- Buckwalter, Wesley and Turri, John (2014), 'Inability and obligation: compelling counterexamples to “ought implies can”', Talk presented at the *Buffalo Experimental Philosophy Conference* (Buffalo, NY). September 20, 2014.
- Buckwalter, Wesley and Turri, John (2015), 'Inability and Obligation in Moral Judgment', *PLOS One*, 10 (8), e0136589. <https://doi.org/10.1371/journal.pone.0136589>.
- Buckwalter, Wesley and Turri, John (Forthcoming), 'Moderate scientism in philosophy', in Jereon de Ridder, Rik Peels, and René van Woudenberg (eds.), *Scientism: Prospects and Problems* (Oxford University Press). Email wesleybuckwalter@gmail.com for penultimate draft.
- Carens, Joseph (2013), *The Ethics of Immigration* (New York: Oxford University Press).

- Chituc, Vladimir, Henne, Paul, Sinnott-Armstrong, Walter, and De Brigard, Felipe (2016), 'Blame, not ability, impacts moral “ought” judgments for impossible actions: Toward an empirical refutation of “ought” implies “can”', *Cognition*, 150, 20-25.
- Cicero, Marcus Tullius and Edmonds, Cyrus R. (1856), *Cicero's three books Of offices, or moral duties : also his Cato major, an essay on old age, Lælius, an essay on friendship, Paradoxes, Scipio's dream, and Letter to Quintus on the duties of a magistrate* (London: H.G. Bohn).
- Copp, David (2008), 'Ought' implies 'can' and the derivation of the principle of alternate possibilities', *Analysis*, 68 (297), 67–75.
- Cruz, Marcio, Foster, James, Quillin, Bryce, and Schellekens, Philip (2015), 'Ending extreme poverty and sharing prosperity : progress and policies. Policy Research Note; PRN/15/03. Washington, D.C. : World Bank Group. <http://documents.worldbank.org/curated/en/801561468198533428/Ending-extreme-poverty-and-sharing-prosperity-progress-and-policies>
- Cushman, Fiery, Durwin, A. J., and Lively, Chaz (2012), 'Revenge without responsibility? Judgments about collective punishment in baseball', *Journal of Experimental Social Psychology*, 48 (5), 1106-10.
- Dahl, Norman O. (1974), 'Ought implies can and deontic logic', *Philosophia*, 4 (4), 485-511.
- Doris, John M. and Plakias, Alexandra (2008), 'How to argue about disagreement: Evaluative diversity and moral realism', in Walter Sinnott-Armstrong (ed.), *Moral*

- psychology, Vol 2: The cognitive science of morality: Intuition and diversity* (Cambridge, MA: MIT Press), 303-31.
- Graham, Jesse, Meindl, Peter, Beall, Erica, Johnson, Kate M., and Zhang, Li (2016), 'Cultural differences in moral judgment and behavior, across and within societies', *Current Opinion in Psychology*, 8, 125-30.
- Graham, P. A. (2011), 'Ought' and Ability', *Philosophical Review*, 120 (3), 337-82.
- Hare, R. M. (1965), *Freedom and Reason* (Oxford: Oxford University Press).
- Henne, Paul, Chituc, Vladimir, De Brigard, Felipe, and Sinnott-Armstrong, Walter (2016), 'An Empirical Refutation of 'Ought' Implies 'Can'', *Analysis*, 76 (3), 283-90.
- Hurley, Paul (2003), 'Fairness and Beneficence', *Ethics*, 113 (4), 841-64.
- Kamm, F. M. (2008), *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (New York: Oxford University Press).
- Kant, Immanuel (1998), *Religion Within the Boundaries of Mere Reason and Other Writings* (Cambridge University Press).
- Knobe, Joshua, Buckwalter, Wesley, Nichols, Shaun, Robbins, Philip, Sarkissian, Hagop, and Sommers, Tamler (2012), 'Experimental Philosophy', *Annual Review of Psychology*, 63 (1), 81-99.
- Kurthy, Miklos and Lawford-Smith, Holly (2015), '[Comment] A brief note on the ambiguity of 'ought'. Reply to Moti Mizrahi's 'Ought, Can and Presupposition: An Experimental Study'', *Method: Analytic Perspectives*, 4 (6), 244-49.

- Kurthy, Miklos, Lawford-Smith, Holly, and Sousa, Paulo (2017), 'Does ought imply can?', *PLOS ONE*, 12(4): e0175206.
<https://doi.org/10.1371/journal.pone.0175206>.
- Littlejohn, Clayton (2012), 'Does 'Ought' Still Imply 'Can'?', *Philosophia*, 40 (4), 821-28.
- Mill, John Stuart (1979), *Utilitarianism*, ed. George Sher (Indianapolis: Hackett).
- Miller, David (2007), *Nationalism and Global Responsibility* (Oxford: Oxford University Press).
- Mizrahi, Moti (2015a), '[Reply] Ought, Can, and Presupposition: A Reply to Kurthy and Lawford-Smith', *Method: Analytic Perspectives*, 4 (6), 250-56.
- Mizrahi, Moti (2015b), 'Ought, Can, and Presupposition: An Experimental Study', *Method: Analytic Perspectives*, 4 (6), 232-43.
- Moore, G. E. (1922), 'The nature of moral philosophy', *Philosophical papers* (Routledge and Kegan Paul).
- O'Neill, Onora (2004), 'Global Justice: Whose Obligations?', in Deen K. Chatterjee (ed.), *The Ethics of Assistance: Morality and the Distant Need* (Cambridge: Cambridge University Press), 242–59.
- O'Neill, Onora (2016), *Justice across Boundaries: Whose Obligations?* (Cambridge: Cambridge University Press).
- Rose, David, Buckwalter, Wesley, and Nichols, Shaun (2017), 'Neuroscientific Prediction and the Intrusion of Intuitive Metaphysics', *Cognitive Science*, 41 (2), 482–502.
- Ryan, Sharon (2003), 'Doxastic compatibilism and the ethics of belief', *Philosophical Studies*, 114 (1-2), 47-79.

- Sarkissian, Hagop (2016), 'Aspects of Folk Morality: Objectivism and Relativism', in Justin Sytsma and Wesley Buckwalter (eds.), *A companion to experimental philosophy* (Hoboken: Wiley), 212-24.
- Sen, Amartya (1999), *Development as Freedom* (Oxford: Oxford University Press).
- Shue, Henry (2004), 'Thickening convergence: Human rights and cultural diversity', in D. Chatterjee (ed.), *The Ethics of Assistance: Morality and the Distant Needy* (Cambridge: Cambridge University Press), 217-41.
- Singer, Peter (1972), 'Famine, affluence, and morality', *Philosophy and Public Affairs*, 1 (3), 229-43.
- Sinnott-Armstrong, Walter (1984), "'Ought' conversationally implies 'can'", *Philosophical Review*, 93 (2), 249-61.
- Stocker, Michael (1971), "'Ought' and 'can'", *Australasian Journal of Philosophy*, 49 (3), 303 – 16.
- Turri, J and Blouw, P (2015), 'Excuse validation: a study in rule-breaking', *Philosophical Studies*, 172 (3), 615-34.
- Turri, John (2013), 'The test of truth: An experimental investigation of the norm of assertion', *Cognition*, 129 (2), 279-91.
- Turri, John (2017a), 'Pathways from Inability to Blamelessness in Moral Judgment', Unpublished manuscript, (University of Waterloo: Philosophical Science Lab).
Email john.turri@gmail.com for draft.
- Turri, John (2017b), 'Compatibilism and Incompatibilism in Social Cognition', *Cognitive Science*, 41 (S3), 403–24.

- UNHCR (2016), 'Forced Displacement in 2015', in United Nations High Commissioner for Refugees (ed.), *Global Trends* (<http://www.unhcr.org/576408cd7.pdf>).
- UNICEF (2014), 'Committing to Child Survival: A Promise Renewed', *Progress Report* (http://files.unicef.org/publications/files/APR_2014_web_15Sept14.pdf).
- UNISDR (2015), '2015 disasters in numbers', (http://www.unisdr.org/files/47804_2015disastertrendsininfographic.pdf: United Nations Office for Disaster Risk Reduction).
- Van Fraassen, Bas C. (1973), 'Values and the heart's command', *Journal of Philosophy*, 70 (1), 5-19.
- Vranas, Peter B. M. (2007), 'I Ought, Therefore I Can', *Philosophical Studies*, 136 (2), 167 - 216.
- Willemsen, Pascale (2016), 'Omissions and expectations: a new approach to the things we failed to do', *Synthese*, 1-28. doi:10.1007/s11229-016-1284-9.