



Inability and Obligation in Intellectual Evaluation

DOI:

[10.1017/epi.2018.49](https://doi.org/10.1017/epi.2018.49)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Buckwalter, W., & Turri, J. (2018). Inability and Obligation in Intellectual Evaluation. *Episteme*. Advance online publication. <https://doi.org/10.1017/epi.2018.49>

Published in:

Episteme

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.



Running Header and short title: INTELLECTUAL EVALUATION

Inability and Obligation in Intellectual Evaluation

Wesley Buckwalter^{abc}

John Turri^d

a University of Manchester, Department of Philosophy, School of Social Sciences,
University of Manchester, Oxford Road, Manchester M13 9PL, UK

b University of Pittsburgh, Department of History and Philosophy of Science, 1101
Cathedral of Learning, 4200 Fifth Avenue, Pittsburgh, PA USA 15260

c Corresponding author, contact email: Wesleybuckwalter@gmail.com

d Philosophy Department and Cognitive Science Program, University of Waterloo, 200
University Avenue West, Waterloo, ON, Canada N2L 3G1

Received: 11-Jul-2018

Revised: 30-Aug-2018

Accepted: 01-Oct-2018

Abstract

If moral responsibilities prescribe how agents ought to behave, are there also intellectual responsibilities prescribing what agents ought to believe? Many theorists have argued that there cannot be intellectual responsibilities because they would require the ability to control whether one believes, whereas it is impossible to control whether one believes. This argument appeals to an “ought implies can” principle for intellectual responsibilities. The present paper tests for the presence of intellectual responsibilities in social cognition. Four experiments show that intellectual responsibilities are attributed to believe things and that these responsibilities can exceed what agents are able to believe. Furthermore, the results show that agents are sometimes considered responsible for failing to form true beliefs on the basis of good evidence, and that this effect does not depend on the seriousness of the consequences for failing to form a belief. These findings clarify when and how responsibilities for belief are attributed, falsify a conceptual entailment between ability and responsibility in the intellectual domain, and emphasize the importance of objective truth in intellectual evaluations.

Keywords: responsibility; belief; truth; evidence; blame; ought implies can; moral psychology

INABILITY AND OBLIGATION IN INTELLECTUAL EVALUATION

1. INTRODUCTION

In his first official White House press briefing, Press Secretary Sean Spicer claimed that the 2017 inauguration of President Trump drew “the largest audience ever to witness an inauguration, period, both in person and around the globe” (White House 2017). This statement was quickly challenged in the popular press when photographs of the National Mall surfaced showing significantly larger crowds during the 2009 inauguration of President Obama. When presented with both photographs however, 41% of Trump voters surveyed in a national poll misidentified the Obama inauguration crowd for the Trump inauguration crowd, and 15% went so far as to claim that the inauguration photo with the objectively sparser crowd actually had more people in it (Schaffner and Luks 2017). These claims are false but does the fact that an agent believes them mean they have done something *wrong*?

Such beliefs are often taken as evidence that we now live in a “post-truth” reality in which emotional appeals guided by personal interests are more important than evidence in shaping what is believed to be true. This rise of post-truth rhetoric connects to a series of fundamental questions about the normative status of belief in folk epistemology, or the suite of concepts and principles that implicitly guide intellectual evaluations in ordinary social cognition. Do we have intellectual obligations to believe certain things? And if so, what factors generate the obligations? For example, do we have an intellectual obligation to believe that one photograph contains more people? If so, is that because the evidence supports the conclusion? One hypothesis, broadly consistent with post-truth relativism, is that there are no intellectual obligations to believe. According to this hypothesis, it is

permissible to believe false things contrary to evidence whenever it suits us to do so. A rival hypothesis is that there are at least some intellectual obligations. A natural adjunct to the alternative hypothesis is that intellectual obligations are due to the evidence we have or what we know to be true. Although beliefs might sometimes be formed by prioritizing personal interests over evidence and objective facts, this is often a mistake.

It is currently unknown if responsibility judgments extend from the moral domain, concerning how one ought to behave, to the intellectual domain, concerning what one ought to believe, although prior research in psychology supports the hypothesis that they do. Several theorists have argued that a fundamental property of moral judgment is to regulate and sustain social relations (Flack and de Waal 2000; Rai and Fiske 2011; Cushman 2013; Malle, Guglielmo, and Monroe 2014). According to the relationship regulation theory of moral judgment, for example, “the core of our moral psychology consists of motives for evaluating and guiding one’s own and others’ judgments and behaviors” to promote social relationships prescribed to sustain groups, and this includes the evaluation of “speech, emotions, attitudes, and intentions” (Rai and Fiske 2011: 59). To the extent that such attitudes promote social cohesion and interpersonal attraction (Lott and Lott 1965; Sachs 1975; Byrne 1997; Friedkin 2004), and partially determine the prescriptions people ultimately accept (e.g. Kohlberg 1984; for discussion see Sripada and Stich 2006), it is reasonable to suppose that responsibility judgments sometimes extend to beliefs.

The existence of intellectual responsibilities remains a foundational question in philosophy, although a host of leading philosophers challenge the hypothesis that responsibilities extend to belief. As plausible as such responsibility attributions might initially seem in psychology, many philosophers have argued that no such responsibilities

exist. According to this line of thought, intellectual responsibilities are ruled out by a pair of plausible theoretical principles. The first principle, known as “doxastic involuntarism,” states that we are unable to control our beliefs (Williams 1973; Alston 1988; Bennett 1990; Pojman 1999; Buckareff 2004; c.f. Turri, Rose, and Buckwalter 2018). The second principle, known as “ought implies can” (“OIC” for short), states that responsibility requires ability (Moore 1922; Streumer 2003; Vranas 2007; Copp 2008). More specifically, this principle states that someone is responsible for doing something only if they are able to do it. Putting these two foundational principles together, then, we cannot have intellectual responsibilities to believe propositions because we are unable to control our beliefs, which is a necessary condition for responsibility (for discussion, see Ryan 2003).

Recent work in experimental cognitive science has demonstrated that the “ought implies can” principle does not accurately describe judgments in some normative domains, such as moral judgment (for a review, see Buckwalter 2017). Several teams of researchers have identified circumstances in which participants attribute moral responsibilities, obligations, and duties to agents for outcomes that they are unable to produce (Buckwalter and Turri 2015; Mizrahi 2015; Chituc, Henne, Sinnott-Armstrong, and De Brigard 2016; Turri 2017a; Turri 2017b). In one set of studies, for example, researchers presented participants with stimuli describing an employer evaluating an employee’s performance (Turri 2017b). Researchers found that participants judged that the employer had a “moral responsibility to give the employee a good evaluation” when she did in fact perform excellently, even if it was “literally impossible” for the employer to do so. These results indicate that inability does not always preclude the attribution of moral responsibility. This

suggests that the “ought implies can” principle may not preclude responsibility attribution in the intellectual domain, but the possibility has not been tested with respect to beliefs.

This possibility raises three central research questions for studying intellectual responsibility, as reflected in ordinary social cognition. First, pre-theoretically, do we recognize intellectual responsibilities to believe propositions at all? Second, assuming that we do recognize at least some such responsibilities, do we attribute intellectual responsibilities to agents who are unable to fulfill them? That is, does commonsense psychology implicitly reject the “ought implies can” principle in the intellectual domain, as it does in the moral? Third, on the basis of what factors do we attribute intellectual responsibilities? In particular, contrary to currently fashionable post-truth rhetoric, do evidence and objective truth-value matter?

This paper presents four experiments advancing our understanding of intellectual responsibility for believing on each of these dimensions. Experiments 1 and 2 demonstrate that we strongly attribute intellectual responsibilities to others, and that we do so when it is both possible and impossible for them to believe a proposition. Experiment 3 demonstrates that intellectual responsibility attributions are significantly predicted by both having good evidence for a proposition and the fact that the proposition is true rather than false. Experiment 4 replicates this effect and demonstrates that evidence and truth continue to predict epistemic responsibility attributions when the practical consequences of failing to believe are both trivial and serious. Overall, the results suggest that folk epistemology recognizes the existence of intellectual responsibilities to believe, that such responsibilities sometimes exceed what agents are able to believe, even when doing so would be

emotionally unpleasant, and that such responsibilities are grounded in both evidence and objective truth-value.

2. EXPERIMENT 1

This experiment tests whether people attribute intellectual obligations to believe a proposition, and whether this occurs even when the agent is unable to fulfill the obligation (i.e. form the relevant belief).

2.1 *Method*

One hundred sixty-five participants were tested (aged 20–69 years, mean age = 35 years; 65 female; 98% reporting English as a native language). Participants were U.S. residents, recruited and tested online using Amazon Mechanical Turk and Qualtrics, and compensated \$0.35 for approximately 2 minutes of their time. Following prior research on ability and obligation judgments, we decided in advance to recruit approximately 40 participants per condition. Repeat participation was prevented within and across experiments. The same recruitment and compensation procedures were used for all experiments reported in this paper.

Participants were randomly assigned to one of four conditions in a 2 (ability: possible, impossible) x 2 (scenario: evaluation, diagnosis) between-subjects design. All participants read a simple story, responded to four test items, then completed a brief demographic questionnaire.

The materials were closely adapted from previous research featuring agents who are able or unable to perform an overt action (Turri 2017b). Instead of focusing on an action, however, the present materials feature agents who are unable to perform a specific

intellectual activity, namely, believe the relevant proposition. More specifically, the ability factor manipulates whether it is “literally impossible” or “possible” for an agent to believe something. To ensure that the results are not due to idiosyncratic features of a single vignette, multiple vignettes were used. The scenario factor manipulated whether the context of belief involved evaluating an employee or diagnosing a pet’s illness. For example, here is the text of the story for the diagnosis scenario (ability manipulation in brackets):

Robert is walking his dog through a nearby park. All the evidence indicates that the dog has rabies. And it does have rabies. In [light/spite] of his deep emotional attachment to the dog, it is [impossible/possible] for him to believe that the dog has rabies. As a matter of brain chemistry, it is [literally impossible/definitely possible] for him to believe that the dog has rabies.

Here is the text of the story for the evaluation scenario (with the ability manipulation again in brackets):

Robert is evaluating the performance of a company employee. All the evidence indicates that the employee performed excellently. And the employee did perform excellently. Given the current condition of Robert’s brain, it is [impossible/possible] for him to believe the employee performed excellently. As a matter of brain chemistry, it is [literally impossible/definitely possible] that he can believe the employee performed excellently.

After reading one of the stories, participants responded to four test items while the story remained at the top of the screen. Following past research, the first item involved a conceptual combination task designed to test the “ought implies can” principle. This item

asked participants to indicate whether they thought the protagonist had an ability and/or intellectual responsibility to believe in a single evaluation. For example, participants who saw the diagnosis scenario were asked to “please choose the option that best applies” from a list of four response options (order rotated randomly):

1. Robert has an intellectual responsibility to believe the dog has rabies, and he is able to do so.
2. Robert has an intellectual responsibility to believe the dog has rabies, but he is not able to do so.
3. Robert does not have an intellectual responsibility to believe the dog has rabies, but he is able to do so.
4. Robert does not have an intellectual responsibility to believe the dog has rabies, and he is not able to do so.

After completing this task, participants were taken to a new screen and presented with the other three test items:

It is possible for Robert to believe that the dog has rabies. (possible)

On a scale of 0% to 100%, how likely is it that Robert will believe that the dog has rabies? (percent)


Robert is blameworthy if he does not believe that that the dog has rabies. (blame)

Responses to possible and blame items were collected on a standard 7-point Likert agreement scale, 1 (“strongly disagree”) – 7 (“strongly agree”), left-to-right across the participant’s screen. Responses to the percent item were collected in a text box that allowed numerical values between 0 and 100 percent. Participants seeing the evaluation scenario

saw similar items where the phrase “the dog has rabies” was replaced with “the employee performed excellently.”

2.2 Results

A preliminary multinomial logistic regression revealed that response to the OIC (“ought implies can”) test item was not significantly predicted by scenario, but it was predicted by ability condition. To explore this effect further, results were collapsed across scenarios and a chi-squared test was used. Collapsing across these scenarios, ability significantly affected responses to the OIC test item, $X^2(3, N = 165) = 110.45, p < .001$, Cramer’s $V = .82$. The “responsible and able” response was selected at rates exceeding chance in possible conditions (83%: binomial test, $p < .001$, test proportion = .25). The “responsible but unable” response was selected at rates exceeding chance in impossible conditions (70%: binomial test, $p < .001$, test proportion = .25, see Fig. 1). A follow-up chi-squared test also showed that the rate of “responsible and able” answers in the possible conditions was greater than the rate of “responsible but unable” answers in the impossible conditions, $X^2(1, N = 165) = 4.24, p = .039$, Cramer’s $V = .16$.



A multivariate analysis of variance showed that responses to possible, percent, and blame items were each affected by ability. There was also an effect of scenario on possible and blame items. No interactions were detected (see Table 1).

Table 1

Experiment 1. Multivariate analysis of variance, Scenario, Wilks' Lambda = .870, $F(3, 159) = 7.95$, $p < .001$, $\eta_p^2 = .130$, Ability, Wilks' Lambda = .344, $F(3, 159) = 100.88$, $p < .001$, $\eta_p^2 = .656$, S*A, Wilks' Lambda = .988, $F(3, 159) = 0.64$, $p = .589$, $\eta_p^2 = .012$.

Factor	Dependent variable											
	Possible				Percent				Blame			
	<i>F</i>	df	<i>p</i>	η_p^2	<i>F</i>	df	<i>p</i>	η_p^2	<i>F</i>	df	<i>p</i>	η_p^2
Scenario	7.68	1, 161	.006	.046	0.26	1, 161	.608	.002	6.67	1, 161	.011	.040
Ability	235.98	1, 161	<.001	.594	280.82	1, 161	<.001	.636	8.11	1, 161	.005	.048
S*A	0.90	1, 161	.345	.006	0.12	1, 161	.726	.001	1.38	1, 161	.242	.009

Follow-up independent samples t-tests showed that mean response for possible, percent, and blame items for the evaluation scenario was significantly lower in the impossible condition than the possible condition. Mean response for possible and percent items for the evaluation scenario was significantly lower in the impossible condition than the possible condition (see Table 2).

Table 2

Experiment 1. Independent samples t-tests by ability.

Measure	Impossible		Possible		<i>t</i>	df	<i>p</i>	<i>MD</i>	95% CI	<i>d</i>
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>						
Evaluation										
Possible	2.31	1.87	6.33	0.97	-12.49	83	<.001	-4.02	-4.66, -3.38	-2.70
Percent	19.79	30.86	87.02	21.56	-11.67	83	<.001	-67.24	-78.70, -55.78	-2.53
Blame	3.86	1.87	5.00	1.79	-2.88	83	.005	-1.14	-1.93, -0.35	-0.30
Diagnosis										
Possible	3.23	2.29	6.78	0.62	-9.45	78	<.001	-3.55	-4.30, -2.80	-2.11
Percent	19.15	26.07	83.63	21.12	-12.15	78	<.001	-64.48	-75.04, -53.91	-2.72
Blame	4.93	1.82	5.40	1.82	-1.17	78	.247	-0.48	-1.29, 0.34	-0.26

Follow-up independent samples t-tests also showed that mean response to possible and blame items for the evaluation scenario was significantly lower in the impossible condition than the possible condition. Mean response to the possible item for the evaluation scenario was significantly lower in the impossible condition than the possible condition (see Table 3).

Table 3

Experiment 1. Independent samples t-tests by scenario.

Measure	<i>t</i>	df	<i>p</i>	<i>MD</i>	95% CI	<i>d</i>
Impossible						
Possible	-1.99	80	.050	-0.92	-1.83, .001	-0.44
Percent	0.10	80	.920	0.64	-11.95, 13.22	0.02
Blame	-2.62	80	.010	-1.07	-1.88, -0.26	-0.58
Possible						
Possible	-2.50	81	.015	-0.45	-0.81, -0.09	-0.55
Percent	0.73	81	.471	3.40	-5.93, 12.73	0.16
Blame	-1.01	81	.316	-0.40	-1.19, 0.39	-0.22

One sample t-tests showed that mean response to the possible item were below the neutral mid-point (=4) in all impossible conditions and above the mid-point in all possible conditions. One sample t-tests also showed that mean response to the blame item was above the mid-point in all conditions except the impossible condition of the evaluation scenario (see Table 4).

Table 4

Experiment 1. One sample t-tests. Test value = 4.

Measure	Impossible					Possible				
	<i>t</i>	df	<i>p</i>	<i>MD</i>	95% CI	<i>t</i>	df	<i>p</i>	<i>MD</i>	95% CI
Evaluation										
Possible	-5.87	41	<.001	-1.69	-2.27, -1.11	15.74	42	<.001	2.33	2.03, 2.62
Blame	-0.50	41	.623	-0.14	-0.73, 0.44	3.67	42	.001	1.00	0.45, 1.55
Diagnosis										
Possible	-2.14	39	.039	-0.78	-1.51, -0.04	28.32	39	<.001	2.78	2.58, 2.97
Blame	3.22	39	.003	0.93	0.34, 1.51	4.86	39	<.001	1.40	0.82, 1.98

2.3 Discussion

The results support the existence of intellectual responsibilities in ordinary social cognition. In two scenarios, the majority of participants indicated that protagonists have an intellectual responsibility to believe certain propositions. The results further indicate participants attribute these responsibilities even when they recognize that the agent is unable to form the relevant belief. Finally, the results indicate that agents are also sometimes considered blameworthy for failing to believe things even though it is impossible for those agents to believe them.

3. EXPERIMENT 2

This experiment tests whether the intellectual responsibility attributions observed in Experiment 1 replicate when using a different method to probe for responsibility judgments.

3.1 *Method*

One hundred and sixty-nine new participants were tested (aged 18–66 years, mean age = 33 years; 55 female; 95% reporting English as a native language). Participants were again randomly assigned to one of four conditions in a 2 (ability: possible, impossible) x 2 (scenario: evaluation, diagnosis) between-subjects design. The materials and testing procedures were identical with those used in Experiment 1, with one exception. In this experiment, the four-option OIC test item was replaced with a single item probing responsibility judgements with a scalar agreement measure. For example, in the diagnosis scenario, the responsibility test item read: “Robert has an intellectual responsibility to believe that the dog has rabies.” Responses were collected on the same 7-point Likert agreement scale, 1 (“strongly disagree”) – 7 (“strongly agree”) used above. After evaluating this item, participants advanced to a new screen and the experiment proceeded identically to Experiment 1.

3.2 *Results*

A preliminary multivariate analysis of variance showed that responses to all dependent variables were affected by assignment to the ability condition. Only the blame item was affected by scenario. No interactions were detected (see Table 5).

Table 5

Experiment 2. Multivariate analysis of variance, Scenario, Wilks' Lambda = .914, $F(4, 162) = 3.82$, $p = .005$, $\eta_p^2 = .086$, Ability, Wilks' Lambda = .322, $F(4, 162) = 85.12$, $p < .001$, $\eta_p^2 = .678$, S*A, Wilks' Lambda = .987, $F(4, 162) = .53$, $p = .715$, $\eta_p^2 = .013$.

Factor	Dependent variable															
	Responsibility				Possible				Percent				Blame			
	F	df	p	η_p^2	F	df	p	η_p^2	F	df	p	η_p^2	F	df	p	η_p^2
Scenario	3.85	1, 165	.052	.023	0.03	1, 165	.855	<.001	1.37	1, 165	.243	.008	8.28	1, 165	.005	.048
Ability	32.99	1, 165	<.001	.167	247.32	1, 165	<.001	.600	318.39	1, 165	<.001	.659	10.80	1, 165	.001	.061
S*A	1.35	1, 165	.246	.008	0.38	1, 165	.541	.002	0.08	1, 165	.776	<.001	0.004	1, 165	.949	<.001

Prior work suggests that blame judgments can bias responsibility judgements when responsibility judgments are collected using scaled responses (Chituc, Henne, Sinnott-Armstrong, and De Brigard 2016). To test for the presence of this bias, a mediation analysis was conducted (Hayes 2013) with responsibility judgments as the outcome variable, assignment to ability condition as the independent variable (coded: 0 = impossible, 1 = possible), and blame judgments as a potential mediator. We found that scalar responsibility judgments were biased in this way: blame judgments mediated the effect of assignment to condition on responsibility judgments, indirect effect = .469 [.18, .81], $z = 3.06$, $p = .002$, $\kappa^2 = .14$.

Given that mediation occurred, a new multivariate analysis was conducted, controlling for blame judgments, to investigate the effects of ability and scenario on the dependent variables. After controlling for blame, the multivariate model was non-significant. This analysis showed that responses to responsibility, possible, and percent items were not significantly affected by scenario. Response to all three variables were affected by ability. There were no interactions (see Table 6).

Table 6

Experiment 1. Multivariate analysis of variance controlling for blame, Scenario, Wilks' Lambda = .960, $F(3, 162) = 2.27, p = .083, \eta_p^2 = .040$, Ability, Wilks' Lambda = .344, $F(3, 162) = 103.20, p < .001, \eta_p^2 = .656$, S*A, Wilks' Lambda = .987, $F(3, 162) = .703, p = .551, \eta_p^2 = .013$, Blame, Wilks' Lambda = .666, $F(3, 162) = 27.08, p < .001, \eta_p^2 = .334$.

Factor	Dependent variable											
	Responsibility				Possible				Percent			
	<i>F</i>	<i>df</i>	<i>p</i>	η_p^2	<i>F</i>	<i>df</i>	<i>p</i>	η_p^2	<i>F</i>	<i>df</i>	<i>p</i>	η_p^2
Scenario	0.20	1, 164	.657	.001	0.41	1, 164	.526	.002	4.98	1, 164	.027	.029
Ability	20.72	1, 164	<.001	.112	221.97	1, 164	<.001	.575	295.35	1, 164	<.001	.643
S*A	2.04	1, 164	.155	.012	0.43	1, 164	.513	.003	0.11	1, 164	.745	.001

Follow-up independent samples t-tests (collapsing across scenario) showed that mean response was significantly lower for all test items in the impossible conditions than in possible conditions (see Fig. 2 and Table 7).

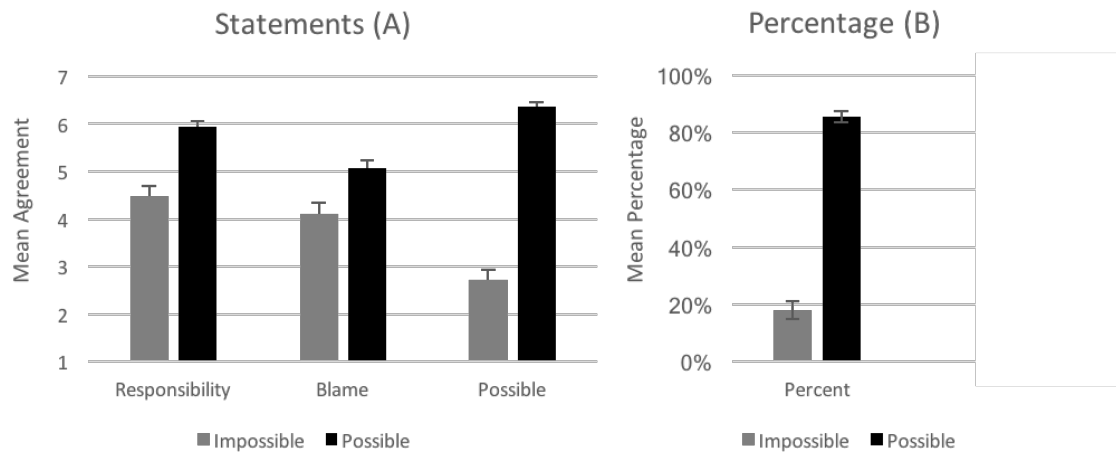


Table 7

Experiment 2. Independent samples t-tests by ability.

Measure	Impossible		Possible		<i>t</i>	<i>df</i>	<i>p</i>	<i>MD</i>	95% CI	<i>d</i>
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>						
Responsible	4.48	2.03	5.93	1.17	-5.67	167	<.001	-1.45	-1.95, -0.94	-0.88
Possible	2.73	1.93	6.36	0.85	-15.80	167	<.001	-3.63	-4.08, -3.17	-2.43
Percent	18.12	29.26	85.56	18.58	-17.87	167	<.001	-67.44	-74.90, -60.00	-2.75
Blame	4.12	2.04	5.06	1.67	-3.29	167	.001	-0.94	-1.51, -0.38	-0.50

Given that there was no effect of scenario after controlling for blame, we also collapsed across scenario to test responsible, possible, and blame items against the neutral midpoint (=4) in both impossible and possible conditions (see Table 8). Mean response to the responsibility item was significantly above the midpoint while mean response to the possible item was significantly below the midpoint in the impossible condition.

Table 8
Experiment 2. One sample t-tests. Test value = 4.

Measure	Impossible					Possible				
	<i>t</i>	df	<i>p</i>	<i>MD</i>	95% CI	<i>t</i>	df	<i>p</i>	<i>MD</i>	95% CI
Responsible	2.19	84	.031	0.48	0.05, 0.92	15.11	83	<.001	1.93	1.67, 2.18
Possible	-6.07	84	<.001	-1.27	-1.69, -0.85	25.56	83	<.001	2.36	2.17, 2.54
Blame	0.53	84	.596	0.12	-0.32, 0.56	5.83	83	<.001	1.06	0.70, 1.42

3.3 Discussion

This experiment replicates the main result of Experiment 1 that the intellectual responsibility to believe propositions is attributed to agents who lack the ability to believe them. Responsibility attributions were significantly affected by ability, but this effect was significantly mediated by blame attributions, and responsibility attribution significantly exceeded the midpoint in the impossible condition. Taken together, the findings from Experiments 1 and 2 suggest that although the method of probing can affect responses, there remains a detectable central tendency to attribute responsibilities despite inability. In other words, the probing method does not shift responses from an anti-OIC pattern to a pro-OIC pattern. Instead, using the scalar probing method depresses the extent to which we observe an anti-OIC pattern.

These results also replicate and expand upon prior findings demonstrating the existence of a bias known as “excuse validation.” According to this bias, the judgment that someone is blameless for accidentally violating a norm can paradoxically lead people to falsely deny that a norm violation occurred in the first place, in order to avoid the appearance that they are unfairly blaming the agent (Turri and Blouw 2015). In Experiment 1, the OIC test item was designed to prevent this bias from occurring, by allowing participants to indicate responsibility and ability in a single evaluation. Because information about inability can be exculpatory, we can, to some extent, avoid triggering excuse validation by allowing participants to evaluate exculpatory information in the same context that they evaluate the responsibility attribution. Alternatively, collecting responsibility and ability judgments in separate test items can promote excuse validation. And the present results are consistent with that happening. Decreased blame judgments for failing to form a belief caused decreased attributions that an intellectual responsibility existed to have been blameworthy for failing to fulfill.

4. EXPERIMENT 3

The prior experiments demonstrate that the attribution of intellectual responsibilities is not necessarily limited by the ability to believe. However, the experiments do not tell us what the intellectual responsibility to believe consists of. This experiment tests whether responsibility attributions are affected by two factors: whether a proposition is supported by none or all of the evidence, and whether the proposition is either objectively false or true. We manipulated evidence and truth-value in the same design in order to detect whether they make independent contributions to responsibility attributions.

4.1 Method

One hundred sixty-two new participants were tested (aged 18–70 years, mean age = 35 years; 78 female; 98% reporting English as a native language). Participants were randomly assigned to one of four conditions in a 2 (evidence: no, all) x 2 (truth-value: true, false) between-subjects design:

Robert is walking his dog through a nearby park. [None/All] of the evidence indicates that the dog has rabies. [But/And] it does [not] have rabies. In light of Robert's deep emotional attachment to the dog, it is impossible for him to believe that the dog has rabies. As a matter of brain chemistry, it is literally impossible for him to believe that the dog has rabies.

After seeing one of four cases, participants were shown items according to a procedure similar to Experiments 1-2. To ensure prior results are further generalizable, we introduced a new OIC test item with only two response options. Participants were asked to “please choose the option that best applies” from the following two options (order rotated randomly):

1. Robert has an intellectual responsibility to believe the dog has rabies, but he is not able to do so.
2. Robert does not have an intellectual responsibility to believe the dog has rabies, and he is not able to do so.

After answering this item participants were taken to a new screen and presented with the following four items:

It is true that the dog has rabies. (truth)

Robert has good evidence that the dog has rabies. (evidence)

It is possible for Robert to believe that the dog has rabies. (possible)

Robert is blameworthy if he does not believe that that the dog has rabies. (blame)

The truth item was answered on a dichotomous “Yes/No” scale. The evidence, possible, and blame items were answered on the same seven-point agreement scale used above.

4.2 Results

A binary logistic regression revealed that both evidence and truth-value significantly predicted responses to the two-option OIC test item. No interaction was detected (see Table 9). The “responsible but unable” response was selected at rates exceeding chance in the true all-evidence condition (76%, binomial test, $p = .001$), but was not significantly selected in the false no-evidence condition (35%, binomial test, $p = .081$), false all-evidence condition (63%, binomial test, $p = .117$), or true no evidence condition (58%, binomial test, $p = .430$, all test proportions = .5).

Table 9

Experiment 3. Logistic regression predicting responses to the OIC test item (2-option). The full model was statistically significant, $X^2(3, N = 162) = 14.64, p < .01$, and explained between 9% and 12% of the variance in response to the test item.

Predictor	B	SE	Wald	df	p	Odds Ratio	95% CI
Evidence	-1.17	.46	6.36	1	.012	0.31	0.13, 0.77
Truth-value	-0.92	.46	4.00	1	.045	0.40	0.16, 0.98
E*T	0.34	.67	0.26	1	.612	1.41	0.38, 5.23
Constant	0.62	.33	3.49	1	.062	1.86	

A binary logistic regression revealed that truth-value (Wald = 25.09, $p < .001$) significantly predicted responses to the truth item. The model was statistically significant, $X^2(3, N = 162) = 108.19, p < .001$, and explained between 49% and 65% of the variance in response to the truth item, with 98/90% answering “yes” in the true no-evidence/all-

evidence conditions (binomial tests, $ps = .001$, test proportions = .5) and 88/73% answering “no” in the false no-evidence/all-evidence conditions (binomial tests, $ps = .001$, test proportions = .5). Evidence was not a significant predictor and no interaction was detected.

A multivariate analysis of variance showed that responses to evidence, possible, and blame items were each affected by evidence. The analysis also demonstrated that evidence and blame items were affected by truth-value. There were no interactions (see Table 10).

Table 10

Experiment 3. Multivariate analysis of variance, Evidence, Wilks' Lambda = .537, $F(3, 156) = 44.83$, $p < .001$, $\eta_p^2 = .463$, Truth, Wilks' Lambda = .872, $F(3, 156) = 7.65$, $p < .001$, $\eta_p^2 = .128$, E*T, Wilks' Lambda = .997, $F(3, 156) = .133$, $p = .940$, $\eta_p^2 = .003$.

Factor	Dependent variable											
	Evidence				Possible				Blame			
	<i>F</i>	<i>df</i>	<i>p</i>	η_p^2	<i>F</i>	<i>df</i>	<i>p</i>	η_p^2	<i>F</i>	<i>df</i>	<i>p</i>	η_p^2
Evidence	133.93	1, 158	<.001	.459	5.20	1, 158	.024	.032	21.29	1, 158	<.001	.119
Truth-value	8.06	1, 158	.005	.049	1.61	1, 158	.207	.010	22.83	1, 158	<.001	.126
E*T	0.02	1, 158	.882	<.001	0.36	1, 158	.548	.002	0.15	1, 158	.702	.001

Follow-up independent samples t-tests showed that mean response for evidence, possible, and blame items was significantly lower in the no-evidence condition than the all-evidence condition when the belief was true. Mean response for evidence and blame items was significantly lower in the no-evidence condition than the all-evidence condition when the belief was false (see Table 11).

Table 11

Experiment 3. Independent samples t-tests by evidence.

Measure	No Evidence		All Evidence		<i>t</i>	<i>df</i>	<i>p</i>	<i>MD</i>	95% CI	<i>d</i>
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>						
False										
Evidence	1.75	1.46	4.83	2.05	-7.77	79	<.001	-3.08	-3.87, -2.29	-1.73

Possible	2.35	2.05	2.85	2.01	-1.12	79	.267	-0.50	-1.40, 0.39	-0.25
Blame	2.65	1.75	3.83	1.72	-3.06	79	.003	-1.18	-1.95, -0.41	-0.68
True										
Evidence	2.48	1.85	5.63	1.41	-8.65	79	<.001	-3.16	-3.89, -2.43	-1.91
Possible	2.55	1.69	3.41	1.87	-2.18	79	.032	-0.87	-1.65, -0.08	-0.48
Blame	3.88	2.07	5.27	1.53	-3.45	79	.001	-1.39	-2.20, -0.59	-0.76

Follow-up independent samples t-tests also showed that mean response to the blame item was significantly lower when the belief was false than when it was true in the no-evidence condition, and that mean response to the evidence and blame items was significantly lower when the belief was false than when it was true in the all-evidence condition (see Table 12).

Table 12

Experiment 3. Independent samples t-tests by truth-value.

Measure	<i>t</i>	df	<i>p</i>	<i>MD</i>	95% CI	<i>d</i>
No Evidence						
Evidence	-1.94	78	.056	-0.73	-1.47, 0.02	-0.44
Possible	-0.48	78	.635	-0.20	-1.04, 0.64	-0.11
Blame	-2.86	78	.005	-1.23	-2.08, -0.37	-0.64
All Evidence						
Evidence	-2.07	80	.041	-0.81	-1.58, -0.03	-0.46
Possible	-1.31	80	.194	-0.56	-1.41, 0.29	-0.29
Blame	-4.00	80	<.001	-1.44	-2.15, -0.72	-0.89

One sample t-tests showed that mean response was significantly below and above the neutral mid-point (=4) for evidence items in no and all evidence conditions, respectively. Tests also showed that mean response to the possible item was significantly below midpoint in both false conditions and the true no-evidence condition, and it was trending in the true all-evidence condition (see Table 13).

Table 13

Experiment 3. One sample t-tests. Test value = 4.

Measure	No Evidence					All Evidence				
	<i>t</i>	df	<i>p</i>	<i>MD</i>	95% CI	<i>t</i>	<i>df</i>	<i>p</i>	<i>MD</i>	95% CI
False										
Evidence	-9.73	39	<.001	-2.25	-2.72, -1.78	2.59	40	.013	0.83	0.18, 1.48
Possible	-5.10	39	<.001	-1.65	-2.30, -1.00	-3.66	40	.001	-1.15	-1.78, -0.51
Blame	-4.89	39	<.001	-1.35	-1.91, -0.79	-0.64	40	.528	-0.17	-0.71, 0.37
True										
Evidence	-5.20	39	<.001	-1.53	-2.12, -0.93	7.42	40	<.001	1.63	1.19, 2.08
Possible	-5.41	39	<.001	-1.45	-1.99, -0.91	-2.00	40	.052	-0.59	-1.18, 0.01
Blame	-0.38	39	.704	-0.13	-0.79, 0.54	5.30	40	<.001	1.27	0.78, 1.75

4.3 Discussion

The results replicate findings from the earlier experiments showing that people attribute intellectual responsibilities, and that they do so even when the agent in question is unable to form the relevant belief. The results also expand upon prior results by identifying two factors, evidence and objective truth, that contribute to the attribution of intellectual responsibilities. The results also replicate prior findings concerning the link between objective truth and the evaluation of evidence (Turri 2015, 2016). In the present study, participants evaluated an agent's evidence as stronger when it supported a claim that was objectively true rather than false.

5. EXPERIMENT 4

This experiment tests whether intellectual obligations are affected by whether the consequences of not believing a proposition are severe.

5.1 Method

One hundred eighty-five new participants were tested (aged 20–74 years, mean age = 36 years; 78 female; 94% reporting English as a native language). Participants were randomly assigned to one of four conditions in a 2 (stakes: low, high) x 2 (epistemic status: bad, good) between-subjects design. The stakes factor manipulated the consequences of failing to form a belief. In the low stakes case, the belief involved diagnosing a pet for dandruff, while the high stakes case involved diagnosing a pet for rabies. The epistemic status factor combined the elements observed to affect responsibility attributions in Experiment 3: the bad epistemic status involved a false proposition supported by no evidence, whereas a good epistemic status involved a true proposition supported by all the evidence. These manipulations are shown in the brackets below:

Robert is walking his dog through a nearby park. [None/All] of the evidence indicates that the dog has [dandruff/rabies]. And it does [not] have [dandruff/rabies]. In light of Robert's deep emotional attachment to the dog, it is impossible for him to believe that the dog has [dandruff/rabies]. As a matter of brain chemistry, it is literally impossible for him to believe that the dog has [dandruff/rabies].

After seeing one of these four cases, participants were presented with the two-option OIC test item used in Experiment 3 asking participants to “please choose the option that best applies” from two options (order rotated randomly):

1. Robert has an intellectual responsibility to believe the dog has [dandruff/rabies], but he is not able to do so.
2. Robert does not have an intellectual responsibility to believe the dog has [dandruff/rabies], and he is not able to do so.

According to the same procedure above, participants were then taken to a new screen and presented with the following agreement items:

It is possible for Robert to believe that the dog has [dandruff/rabies]. (possible)

If Robert does not believe that the dog has [dandruff/rabies], there could be serious consequences. (consequences)

Robert is blameworthy if he does not believe that that the dog has [dandruff/rabies]. (blame)

Responses were collected on the same seven point scales used in prior experiments.

5.2 Results

A binary logistic regression revealed that the two-option OIC test item was predicted by epistemic status. It was not predicted by stakes and no interaction was detected (see Table 14). To explore this effect further, results were collapsed across stakes and a chi-squared test was used. Collapsing across stakes, epistemic status significantly affected responses to the OIC test item, $X^2(1, N = 185) = 40.97, p < .001$, Cramer's $V = .47$. The “responsible but unable” response was selected at rates exceeding chance in the good epistemic condition (72%: binomial test, $p < .001$, test proportion = .50). The “not responsible and unable” response was selected at rates exceeding chance in the bad epistemic condition (75%: binomial test, $p < .001$, test proportion = .50).

Table 14

Experiment 4. Logistic regression predicting responses to the OIC test item (2-option). The full model was statistically significant, $X^2(3, N = 185) = 50.63, p < .001$, and explained between 24% and 32% variance in response to the test question.

Predictor	B	SE	Wald	df	<i>p</i>	Odds Ratio	95% CI
Stakes	-0.65	0.49	1.73	1	.189	0.52	0.20, 1.37
Epistemic Status	-1.90	0.47	16.11	1	<.001	0.15	0.06, 0.38
S*ES	-0.56	0.71	0.64	1	.425	0.57	0.14, 2.27
Constant	1.44	0.37	15.10	1	<.001	4.22	

A multivariate analysis of variance showed that responses to the consequences and blame items were each affected by stakes and epistemic status. There were no interactions (see Table 15).

Table 15

Experiment 4. Multivariate analysis of variance, Stakes, Wilks' Lambda = .640, $F(3, 179) = 33.62$, $p < .001$, $\eta^2 = .360$, Epistemic, Wilks' Lambda = .721, $F(3, 179) = 23.13$, $p < .001$, $\eta^2 = .279$, S*E, Wilks' Lambda = .976, $F(3, 179) = 1.47$, $p = .225$, $\eta^2 = .024$.

Factor	Dependent variable											
	Possible				Consequences				Blame			
	<i>F</i>	<i>df</i>	<i>p</i>	η^2	<i>F</i>	<i>df</i>	<i>p</i>	η^2	<i>F</i>	<i>df</i>	<i>p</i>	η^2
Stakes	0.01	1, 181	.938	<.001	100.01	1, 181	<.001	.356	9.32	1, 181	.003	.049
Epistemic Status	3.04	1, 181	.083	.017	51.34	1, 181	<.001	.221	45.65	1, 181	<.001	.201
S*ES	0.82	1, 181	.365	.005	3.68	1, 181	.057	.020	0.20	1, 181	.659	.001

Follow-up independent samples t-tests showed that mean response for consequences and blame items was significantly lower in the bad epistemic status condition than the good, but that possibility was not affected by status (see Table 16).

Table 16

Experiment 4. Evidence independent samples t-tests by epistemic status.

Measure	Bad Epistemic		Good Epistemic		<i>t</i>	<i>df</i>	<i>p</i>	<i>MD</i>	95% CI	<i>d</i>
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>						
Low Stakes										
Possible	2.79	1.79	3.02	2.08	-0.59	94	.558	-0.23	-1.02, 0.55	-0.12
Consequences	2.60	1.41	3.80	1.61	-3.88	94	<.001	-1.20	-1.81, -0.59	-0.79
Blame	2.72	1.46	4.31	1.79	-4.74	94	<.001	-1.58	-2.25, -0.92	-0.97
High Stakes										
Possible	2.56	1.71	3.30	1.97	-1.89	87	.062	-0.74	-1.52, 0.04	-0.40
Consequences	4.44	2.10	6.52	0.82	-6.14	87	<.001	-2.08	-2.75, -1.41	-1.31
Blame	3.38	1.95	5.18	1.57	-4.80	87	<.001	-1.80	-2.55, -1.06	-1.02

Follow-up independent samples t-tests showed that mean response for consequences and blame items was significantly lower in the low stakes condition than the high, but that possibility was not affected by stakes (see Table 17).

Table 17
Experiment 4. Independent samples t-tests by stakes.

Measure	<i>t</i>	df	<i>p</i>	<i>MD</i>	95% CI	<i>d</i>
Bad Epistemic						
Possible	0.63	90	.528	0.23	-0.50, 0.96	0.13
Consequences	-4.99	90	<.001	-1.85	-2.59, -1.11	-1.04
Blame	-1.83	90	.070	-0.65	-1.36, 0.06	-0.38
Good Epistemic						
Possible	-0.65	91	.515	-0.28	-1.11, 0.56	-0.14
Consequences	-10.13	91	<.001	-2.73	-3.26, -2.19	-2.14
Blame	-2.49	91	.015	-0.88	-1.57, -0.18	-0.52

One sample t-tests showed that mean response was significantly below the neutral mid-point (=4) for all possible items (see Table 18).

Table 18
Experiment 4. One sample t-tests. Test value = 4.

Measure	Bad Epistemic					Good Epistemic				
	<i>t</i>	df	<i>p</i>	<i>MD</i>	95% CI	<i>t</i>	df	<i>p</i>	<i>MD</i>	95% CI
Low Stakes										
Possible	-4.64	46	<.001	-1.21	-1.74, -0.69	-3.30	48	.002	-0.98	-1.58, -0.38
Consequences	-6.83	46	<.001	-1.40	-1.82, -0.99	-0.89	48	.378	-0.20	-0.67, 0.26
Blame	-6.01	46	<.001	-1.28	-1.7, -0.85	1.20	48	.238	0.31	-0.21, 0.82
High Stakes										
Possible	-5.66	44	<.001	-1.44	-1.96, -0.93	-2.37	43	.022	-0.71	-1.30, -0.11
Consequences	1.42	44	.162	0.44	-0.18, 1.07	20.39	43	<.001	2.52	2.27, 2.77
Blame	-2.15	44	.037	-0.62	-1.21, -0.04	4.98	43	<.001	1.18	0.70, 1.66

5.3 *Discussion*

This experiment replicates the finding that intellectual responsibilities to believe propositions are attributed even to agents unable to fulfill them. Responsibility attribution was significantly predicted by epistemic status — namely, having good evidence for a true proposition, compared to having no evidence for a false proposition. Responsibility attribution was not significantly predicted by stakes, or how costly it would be to fail to form the belief.

6. GENERAL DISCUSSION

The results of four experiments support the existence of intellectual responsibilities in social cognition. First, people clearly attribute intellectual responsibilities to believe propositions. This finding was replicated using three different probing methods, including one scalar measure (Experiment 2) and two different categorical measures (Experiment 1, 3-4). The experiments also revealed two factors contributing to the attribution of intellectual responsibilities: the evidence favoring a proposition, and the proposition's objective truth-value (Experiment 3-4). Finally, the experiments demonstrate that the attribution of these responsibilities is quite strong, sometimes persisting even when the agent in question is unable to form the relevant belief (Experiments 1-4).

The findings also have theoretical significance. In addition to advancing our understanding of how intellectual responsibilities are attributed, these results also cast further doubt on the alleged commitment to “ought implies can” principles in ordinary social cognition. Prior results have demonstrated that ordinary moral judgments violate this principle. The present results extend this finding from the moral to the intellectual domain,

undermining the claim that ordinary social cognition implicitly accepts an “ought implies can” for intellectual responsibilities (e.g. Boghossian 2006: 23). Of course, these findings do not demonstrate that intellectual obligations are never influenced by considerations about what one is able to do. By changing the materials or stimuli, future researchers might profitably explore cases in which inability does intuitively preclude responsibility. The present results demonstrate one set of circumstances in which inability does not do this, which falsifies a general conceptual entailment between ability and responsibility in the intellectual domain.

The results also offer a reason for optimism concerning our political dialog and folk-epistemological assessments of belief. To the extent that people attribute responsibility in these cases, the results suggest that they do not endorse post-truth relativism about objective truth. Instead, people recognize that failing to form beliefs on the basis of true evidence violates an intellectual responsibility of belief. Such responsibilities are so strong, in fact, that they sometimes extend beyond our ability to fulfill them. In the present set of experiments, it appears that both objective truth-value and the possession of evidence both affected responses, and that when combined they were enough to trigger intellectual responsibility attribution. However, further research might profitably explore whether these factors ever trigger intellectual responsibility attributions in isolation, such as a responsibility to believe objectively true propositions without evidence.

Future research could improve our understanding of intellectual responsibility by expanding on the present results in several ways. For one, the scenarios tested responsibility attributions in situations where it was stipulated that claims were true and

were clearly supported by good evidence. For another, these situations involved nonpartisan beliefs that were unassociated with any particular political cause or social group. These features of the materials have the benefit of isolating and identifying the effect of evidence and truth in general intellectual evaluations. However, it is often unclear whether claims are true or well supported and this may be compounded by contentious political content. Further research is needed to test the role of truth or evidence judgments on intellectual responsibility attributions in these more specific circumstances.

Acknowledgements —For helpful feedback, we thank Carolyn Buckwalter and Angelo Turri. This research was supported by the Social Sciences and Humanities Research Council of Canada and the Canada Research Chairs Program.

REFERENCES

- Alston, William P. (1988), 'The deontological conception of epistemic justification', *Philosophical Perspectives*, 2, 257-99.
- Bennett, Jonathan (1990), 'Why is Belief Involuntary?', *Analysis*, 50 (2), 87-107.
- Boghossian, Paul (2006), *Fear of knowledge: against relativism and constructivism* (Oxford: Oxford University Press).
- Buckareff, Andrei A. (2004), 'Acceptance and deciding to believe', *Journal of Philosophical Research*, 29, 173-90.
- Buckwalter, Wesley (2017), 'Ability, Responsibility, and Global Justice', *Journal of Indian Council of Philosophical Research*, 34(3), 577-590.
- Buckwalter, Wesley and Turri, John (2015), 'Inability and Obligation in Moral Judgment', *PLOS One*, 10 (8), <https://doi.org/10.1371/journal.pone.0136589>.

- Byrne, Donn (1997), 'An Overview (and Underview) of Research and Theory within the Attraction Paradigm', *Journal of Social and Personal Relationships*, 14 (3), 417-31.
- 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.
- Copp, David (2008), 'Ought' implies 'can' and the derivation of the principle of alternate possibilities', *Analysis*, 68 (297), 67–75.
- Cushman, Fiery (2013), 'The role of learning in punishment, prosociality, and human uniqueness', in Kim Sterelny, et al. (eds.), *Cooperation and Its Evolution* (Cambridge: MIT Press), 333-72.
- Flack, J. C. and de Waal, F. B. M. (2000), 'Any animal whatever'. Darwinian building blocks of morality in monkeys and apes', *Journal of Consciousness Studies*, 7 (1-2), 1-29.
- Friedkin, Noah E. (2004), 'Social Cohesion', *Annual Review of Sociology*, 30 (1), 409-25.
- Hayes, Andrew F. (2013), *Introduction to Mediation, Moderation, and Conditional Process Analysis: A Regression-Based Approach* (New York, NY: Guilford Press).
- Kohlberg, Lawrence (1984), *The Psychology of Moral Development: The Nature and Validity of Moral Stages* (San Francisco: Harper & Row).
- Lott, A. J. and Lott, B. E. (1965), 'Group cohesiveness as interpersonal attraction: a review of relationships with antecedent and consequent variables', *Psychol Bull*, 64 (4), 259-309.

- Malle, Bertram F., Guglielmo, Steve, and Monroe, Andrew E. (2014), 'A Theory of Blame', *Psychological Inquiry*, 25 (2), 147-86.
- Mizrahi, Moti (2015), '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).
- Pojman, L. P. (1999), 'Believing, Willing, and the Ethics of Belief', *The Theory of Knowledge: Classical and Contemporary Readings* (Belmont, CA: Wadsworth), 525-43.
- Rai, T. S. and Fiske, A. P. (2011), 'Moral psychology is relationship regulation: moral motives for unity, hierarchy, equality, and proportionality', *Psychol Rev*, 118 (1), 57-75.
- Ryan, Sharon (2003), 'Doxastic compatibilism and the ethics of belief', *Philosophical Studies*, 114 (1-2), 47-79.
- Sachs, Donald H. (1975), 'Belief similarity and attitude similarity as determinants of interpersonal attraction', *Journal of Research in Personality*, 9 (1), 57-65.
- Schaffner, Brian and Luks, Samantha (2017), 'This is what Trump voters said when asked to compare his inauguration crowd with Obama's', <https://www.washingtonpost.com/news/monkey-cage/wp/2017/01/25/we-asked-people-which-inauguration-crowd-was-bigger-heres-what-they-said/?utm_term=.c86d04ae6a2a>, accessed August 22.

- Sripada, Chandra and Stich, Stephen (2006), 'A Framework for the Psychology of Norms', in Peter Carruthers, Stephen Laurence, and Stephen P. Stich (eds.), *The Innate Mind, Volume 2: Culture and Cognition* (Oxford: Oxford University Press).
- Streumer, Bart (2003), 'Does 'ought' conversationally implicate 'can'?', *European Journal of Philosophy*, 11 (2), 219–28.
- Turri, John and Blouw, Peter (2015), 'Excuse validation: a study in rule-breaking', *Philosophical Studies*, 172 (3), 615-34.
- Turri, J. (2017a), 'How "ought" exceeds but implies "can": Description and encouragement in moral judgment', *Cognition*, 168, 267-75.
- Turri, John (2015), 'Evidence of factive norms of belief and decision', *Synthese*, 192 (12), 4009-30.
- Turri, John (2016), 'The Radicalism of Truth-insensitive Epistemology: Truth's Profound Effect on the Evaluation of Belief', *Philosophy and Phenomenological Research*, 93 (2), 348-67.
- Turri, John (2017b), 'Compatibilism and Incompatibilism in Social Cognition', *Cognitive Science*, 41 (S3), 403–24.
- Turri, John, Rose, D, and Buckwalter, W (2018), 'Choosing and refusing: doxastic voluntarism and folk psychology', *Philosophical Studies*, 175, 2507-2537.
- Vranas, Peter B. M. (2007), 'I Ought, Therefore I Can', *Philosophical Studies*, 136 (2), 167 - 216.
- White House, Office of the Press Secretary (2017), 'Press Briefing by Press Secretary Sean Spicer', (January 21, 2017).

Williams, Bernard (1973), 'Deciding to believe', *Problems of the Self* (Cambridge: Cambridge University Press), 136-51.

FIGURE CAPTIONS

Fig. 1. Experiment 1. Mean percentage choosing each OIC test item option grouped by ability.

Fig. 2. Experiment 2. Panel A: mean response to test statements about whether an agent had an intellectual responsibility, deserves blame, and had a possibility to believe. Scales ran 1 (SD)–7 (SA). Panel B: mean estimate of the percentage chance that the agent would form the belief. Scales ran 0%–100%. All error bars represent +/- 1 SEM.