

Does Non-Moral Ignorance Exculpate? Situational Awareness and Attributions of Blame and Forgiveness

Abstract: In this paper, we set out to test empirically an idea that many philosophers find intuitive, namely, that non-moral ignorance can exculpate. Many philosophers find it intuitive that moral agents are responsible only if they know the particular facts surrounding their action (or inaction). Our results show that whether moral agents are aware of the facts surrounding their (in)action does have an effect on people's attributions of blame, regardless of the consequences or side effects of the agent's actions. In general, it was more likely that a situationally aware agent will be blamed for failing to perform the obligatory action than a situationally unaware agent. We also tested attributions of forgiveness in addition to attributions of blame. In general, it was less likely that a situationally aware agent will be forgiven for failing to perform the obligatory action than a situationally unaware agent. When the agent is situationally unaware, it is more likely that the agent will be forgiven than blamed. We argue that these results provide some empirical support for the hypothesis that there is something intuitive about the idea that non-moral ignorance can exculpate.

Keywords: awareness; blame; blameworthiness; forgiveness; ignorance; moral cognition; moral psychology; moral responsibility

1. Introduction

Many philosophers find it intuitive that moral responsibility has an epistemic dimension.¹ An epistemic condition on moral responsibility, as Fischer and Ravizza (1998, p. 13) put it, is supposed to “*capture the intuition* that an agent is responsible only if he both knows the particular facts surrounding his action, and acts with the proper sorts of beliefs and intentions” (emphasis added). Accordingly, arguments for such an epistemic condition on moral responsibility usually take the form of appeals to intuitions elicited from hypothetical cases (also known as “the method of cases”).² For instance, Fischer and Ravizza (1998, p. 12) consider the case of, “the person who backs his car out of his garage unaware that a tiny kitten is snoozing beneath the rear tire,” and say that “it would be odd to judge that the driver is morally responsible for the kitten's untimely death” (cf. Mele 2011).³

Similarly, by considering hypothetical cases designed to elicit intuitions about epistemic conditions for being blameworthy, Ginet (2000, pp. 276-277) argues for the following epistemic condition on blameworthiness:

either (i) S knew* [i.e., “actively believed” (Ginet 2000, p. 270)] at t₀ that her acting (or omitting to act) in way W would or might bring about (or, in the case of omission, prevent) a harmful act (or omission) of the sort S subsequently committed at t₁, or (ii) S did not at t₀ know* this but there is a sequence of one or more acts (or omissions) that ends with the act (or omission) at t₀ and is such that (a) each member before this last member benights the subsequent member, (b) the first (earliest) member of the sequence was not a benighted act (or omission)—at the time of it S knew* that it would or might lead to the sort of harm it in fact led to, namely, the benighted act (or omission) that is the

¹ For an overview of the literature on the epistemic condition on moral responsibility, see Wieland (2017).

² On the method of cases, see Nagel (2012), Mizrahi (2014), and Baz (2016).

³ According to Wieland (2017), “As many say since Strawson (1962), S is morally blameworthy for X (whether it is an action, omission, consequence, etc.) iff on account of X, S is a proper candidate for the reactive attitudes, such as resentment, indignation, and condemnation.”

next member of the sequence, and (c) at the time of each benighted act (or omission) in the sequence S should have known* (was blameworthy for not knowing*) that it would or might lead to the sort of benighted act (or omission) that it in fact led to in the next member of the sequence.⁴

Ginet's argument for this epistemic condition on moral responsibility is an appeal to intuitions elicited from hypothetical cases like the following:

Simon enters the hotel room he has just checked into and flips what appears to be, and what he takes to be, an ordinary light switch, but, to his surprise and consternation, the flipping of the switch sets off a loud fire alarm (Ginet 2000, p. 269).

To Ginet (2000, p. 269), "*It seems that*, because [Simon] did not know that his flipping the switch would have this unfortunate consequence, it would be wrong to feel indignant with him for bringing about that consequence" (emphasis added).⁵ Rosen (2008, p. 591) calls the claim that "[w]hether a person is morally responsible for a bad act often turns decisively on what he knew and when he knew it" a "platitude."⁶

As Tognazzini (2010) observes, however, "With a few notable exceptions, most who write on moral responsibility acknowledge that it has an epistemic dimension -- in a paragraph, if they are given to wordiness -- before quickly pressing on to sexier questions about control and free will."⁷ For example, Harman (2011, p. 444) finds it natural to think "that ignorance of non-

⁴ Cf. Sher's (2009, p. 143) "full epistemic condition," which is also stated in terms of awareness and lack of awareness. For a discussion of the epistemic requirements for moral responsibility in the context of criminal law, see Oshana (2013).

⁵ See also Fischer and Tognazzini (2009). Cf. Mele (2011).

⁶ In addition to an epistemic condition, there may be other conditions for blameworthiness (and for moral responsibility in general, see, e.g., Cova 2017). For example, some have argued for the principle known as "Ought Implies Can" (OIC) on the grounds that to blame moral agents who are unable to do what they ought to do would be unfair. See Fischer (2003). Cf. Graham (2011). In this paper, we are interested in the epistemic dimension of moral responsibility; in particular, non-moral ignorance. For more on OIC, see Mizrahi (2009) and (2012).

⁷ See, e.g., Haji (1997). On blameworthiness and control, see Brekke Carlsson (2017).

moral matters can exculpate.” She briefly mentions one hypothetical case that is supposed to establish that before quickly moving on to discuss whether moral ignorance can also exculpate:

Suppose Anne spoons some cyanide into Bill’s coffee, but she does not know it is cyanide: she believes that she is spooning sugar. *It seems that Anne is blameworthy for poisoning Bill only if she is blameworthy for her ignorance.* If we stipulate that Anne is not blameworthy for her ignorance, then it seems Anne is not blameworthy for poisoning Bill (emphasis added).⁸

We find the epistemic dimension of moral responsibility quite sexy and worthy of experimental attention. So we propose that an empirical approach can shed new light on the epistemic dimension of moral responsibility. More explicitly, we designed our study to test experimentally “the intuition that an agent is responsible only if he [...] knows the particular facts surrounding his action” (Fischer and Ravizza 1998, p. 13). In other words, we would like to find out whether what might be called “situational awareness,” i.e., a moral agent’s awareness or lack of awareness of the particular facts surrounding his or her (in)action, has any effect on attributions of blameworthiness and forgiveness. If this epistemic condition on moral responsibility is indeed *intuitive*, as the aforementioned philosophers claim, then we would expect people to withhold blame from moral agents who are situationally unaware, i.e., unaware of the particular facts surrounding their action (or inaction). On the other hand, if we find that people attribute blame to situationally unaware agents, then that would suggest that it does not *seem* to people that situationally unaware agents are unworthy of blame, and hence that people do not find it *intuitive* “that an agent is responsible only if he [...] knows the particular facts surrounding his action” (Fischer and Ravizza 1998, p. 13). This result, if obtained, would suggest that there is something intuitive about the idea of a non-moral, epistemic dimension for moral responsibility;

⁸ Cf. Rosen (2004). See also Wieland (2015).

specifically, an epistemic dimension having to do with moral agents' situational awareness (or lack thereof).

Before we discuss our experimental study in detail, a few important clarifications are in order. First, following the work of Knobe (see, e.g., 2003a, 2003b, 2004a, 2004b), on what is now known as the “Knobe Effect” (namely, that the perceived goodness or badness of an action's side effects affects people's attributions of intentionality), there has been a lot of experimental work on attributions of blame, intentionality, and intentional action (see, e.g., Cushman 2008).⁹ By contrast, our study is *not* designed to test attributions of blame to agents who bring about harm intentionally or unintentionally (as in Cole and Bengson 2009, for example). In fact, our vignettes do not mention the consequences of the agent's actions or side effects at all. What we are interested in is *not* whether intentionality affects attributions of blame, but rather whether *facts about situational awareness* (or lack of situational awareness) affect attributions of blame as well as attributions of forgiveness. There has been some experimental work on *moral* ignorance due to upbringing and attributions of blame (see, e.g., Faraci and Shoemaker 2014). However, to the best of our knowledge, there are no existing experimental studies on non-moral or factual ignorance, specifically, *situational awareness (or lack of awareness)*, that directly test the hypothesis that it is intuitive an agent is morally responsible only if she knows the particular facts surrounding her (in)action; that is, “that ignorance of non-moral matters can exculpate” (Harman 2011, p. 444).

Second, it is important to carefully distinguish between moral agents whose actions have unforeseeable consequences and moral agents who are situationally aware (or situationally unaware), i.e., aware or unaware of the particular facts surrounding their (in)action. In this study,

⁹ For an overview of this literature, see Feltz (2007) and Cova (2016). For a critical discussion, see Nadelhoffer (2004) and Sauer (2014).

we are interested in the latter, not the former. As for the former, philosophers sometimes talk about a “foreseeability condition” in the context of a so-called “tracing principle,” according to which

When one acts from a reasons-responsive mechanism at time $T1$, and one can reasonably be expected to know that so acting will (or may) lead to acting from an unresponsive mechanism at some later time $T2$, one can be held responsible for so acting at $T2$ (Fischer and Ravizza 1998, p. 50).

Notice that, according to the “tracing principle,” what must be foreseeable is *not* the consequences of one’s action. Rather, what must be foreseeable is one’s *incapacitation*. For example:

The intoxicated wrongdoer is blameworthy despite his incapacitation precisely because he is responsible for becoming incapacitated. We hold him responsible for his intoxicated wrongdoing by tracing back to his responsibility for becoming intoxicated (Agule 2016, p. 1).

Similarly, Benchimol (2011, p. 102) offers the following case: “a lifeguard on duty (call her “Abby”) falls asleep on the job, during which time an inexperienced swimmer is pulled away from the shore by a strong undertow, and Abby doesn’t assist him as she should,” and says that “Abby is blameworthy only if her unintentional omission is either an explicitly foreseen or reasonably foreseeable consequence of some prior choice that reflected ill will or disregard” (Benchimol 2011, p. 107).¹⁰ There is some empirical evidence suggesting that the *foreseeability of consequences* affects attributions of blame (Lagnado and Channon 2008). To the best of our knowledge, however, there is no empirical evidence pertaining to the foreseeability of *incapacitation* or *omission*. Be that as it may, our study is *not* designed to test the effects of

¹⁰ For arguments against the so-called “tracing principle,” see Khoury (2012).

foreseeability (of consequences, incapacitation, or omission). Rather, our study is designed to test “the intuition that an agent is responsible only if he [...] knows the particular facts surrounding his action” (Fischer and Ravizza 1998, p. 13). We are interested in how facts about moral agents’ situational awareness (or lack of situational awareness) affect attributions of blame and forgiveness.¹¹

Third, our study is also *not* about what Sher (2009, p. 23) calls “unwitting wrongdoers and foolish agents.” Such agents are unaware of the *wrongness* of their actions, whereas our agents are unaware of the facts surrounding their (in)action, which is what we call “situational awareness.” So our cases are different from Sher’s (2009, pp. 23-40) insofar as they are about agents being unaware of *circumstances*, not *moral wrongness*, and they are designed to test attributions of blame (as well as forgiveness) regardless of the consequences or side effects of the agents’ (in)action. In other words, the sort of ignorance we are interested in is factual ignorance, specifically, situational ignorance, not moral ignorance, i.e., when “the agent may not have realized that her action was wrong” (Montminy 2016, p. 56).¹² In our vignettes, the agent is not morally ignorant, i.e., the agent is not ignorant of moral facts. Rather, the agent is situationally ignorant, i.e., ignorant of the facts surrounding his/her (in)action, in particular, that that a person using crutches just boarded the train and there are no available seats.¹³ As we will see in Section 2c, our participants judge that the agents are morally obligated to perform an action despite being situationally unaware.

¹¹ According to Zaibert (2009, p. 388), the “standard view” is that “to forgive is, explicitly, to overcome resentment (even if this overcoming is done for a moral reason,” but Zaibert (2009, p. 368) argues that “to forgive is to deliberately refuse to punish.” Cf. Warmke (2011).

¹² See also Wieland (2015) for the contrast between moral ignorance and factual ignorance.

¹³ In Luke 23:34, Jesus says “Father, forgive them; for they know not what they do.” Presumably, Jesus is talking about moral ignorance here; that is, they do not know that what they are doing is morally wrong, which is different from being ignorant of non-moral facts. For more on moral ignorance, see Guerrero (2007).

In the next section, then, we discuss the methods and results of our experimental study on non-moral or factual ignorance (specifically, situational awareness) and moral responsibility (specifically, blame and forgiveness). Our results show that whether moral agents (the subjects of blame or forgiveness) are aware of the facts surrounding their (in)action (e.g., giving up a seat on a train for a person using crutches) does have an effect on people's attributions of blame. When the agent is situationally aware, people are more inclined to blame the agent for failing to perform the obligatory action. We also tested attributions of forgiveness in addition to attributions of blame. When the agent is situationally aware, people are less inclined to forgive the agent for failing to perform the obligatory action. When the agent is situationally unaware, people are more inclined to forgive than blame. We will argue that these results provide some empirical support for the hypothesis that there is something intuitive or "natural" about the "thought that ignorance of non-moral matters can exculpate" (Harman 2011, p. 444).

2. Study

2a. Methods

Our experimental study was designed to test people's attributions of obligation (ought), blame, and forgiveness in cases where moral agents are either aware or unaware of the facts surrounding their (in)action. In particular, we wanted to see whether moral agents (the subject to whom obligation, blame, or forgiveness can be attributed) are situationally aware or unaware has any effect on people's attributions of obligation, blame, and forgiveness. To this end, we manipulated the following variables:

Circumstances (Situationally Aware/Situationally Unaware): We gave a portion of participants vignettes in which the agent is aware of the facts surrounding his/her

(in)action (the agent is *awake*), and other participants vignettes in which the agent is unaware of the facts surrounding his/her (in)action (the agent is *asleep*).

Agent (Alex/You): We gave a portion of participants vignettes in which they are observers (the agent is “Alex”) and other participants vignettes in which they are actors (the agent is “You”).¹⁴

Participants for this study were recruited through Amazon Mechanical Turk and were tested through Qualtrics, an online survey platform. They were compensated \$0.15 for approximately two minutes of their time. A total of 1084 participants were recruited for this study. Of this sample, three participants were excluded because they did not consent to the informed agreement and 11 participants were excluded for not completing the questionnaire. This resulted in a sample of 1070 total participants.¹⁵

After they were given written informed consent, participants were randomly given one vignette in a between-subjects experimental design that read as follows:

You are [/Alex is] on a train. A person with a broken leg who uses crutches to walk boards the train. You remain [/Alex remains] seated even though there are no available seats on the train.

In the vignettes in which the agent is situationally unaware (the agent is asleep, and so is unaware of the facts that a person using crutches just boarded the train and that there are no available seats), the second sentence read as follows:

¹⁴ We chose a gender-neutral name, like “Alex,” in an attempt to avoid as much as possible any gender-related influences on people’s judgments.

¹⁵ Although demographics had no statistical effects, it is worth noting that our sample was extremely diverse, with participants (ages 18-65) claiming at least three different genders (624 men, 433 women, and 6 other), 17 different ethnic backgrounds, and 25 distinct religious affiliations.

You are [/Alex is] sleeping and remain[/s] seated even though there are no available seats on the train.

After reading one of the vignettes, participants were randomly assigned to one of the following statements and asked to indicate their agreement or disagreement with them on a standard Likert scale from 1 (“Strongly Disagree”) to 7 (“Strongly Agree”):

1. You [/Alex] ought to give up the seat.
2. You are [/Alex is] worthy of blame for not giving up the seat.
3. You are [/Alex is] worthy of forgiveness for not giving up the seat.

We have chosen a between-subjects experimental design, randomly assigning each participant to a vignette with one of the above statements, for the following reasons. First, we wanted to avoid interference between the conditions, and so make our results more generalizable. Second, since we wanted to know whether participants blame or forgive equally for the two vignettes, a between-subjects design is preferable, given that a within-subjects design could only tell us when participants blame (or forgive) with respect to the vignette they originally saw. This design could then introduce cognitive biases and ordering effects. As far as cognitive biases are concerned, participants may feel that they must pick a different value for the different questions asked, which, in turn, would force a lower blame assignment for vignette in which the agent is unaware. But these results would not represent the blame level participants really wanted to assign. Instead, such results would show the blame participants were willing to assign given what they had assigned earlier. In such a case, we would then see a decrease in the level of blame assigned, but it would *not* be the level of blame assigned in general. As for ordering effects, a within-subjects experimental design could introduce an ordering effect that can easily be avoided, given that, in such an experimental design, it is going to matter which vignette participants read first.

For these reasons, then, we have chosen a between-subjects experimental design. After indicating their agreement or disagreement with one of the aforementioned statements, 1-3, participants were asked to complete a brief demographics questionnaire.

2b. Results

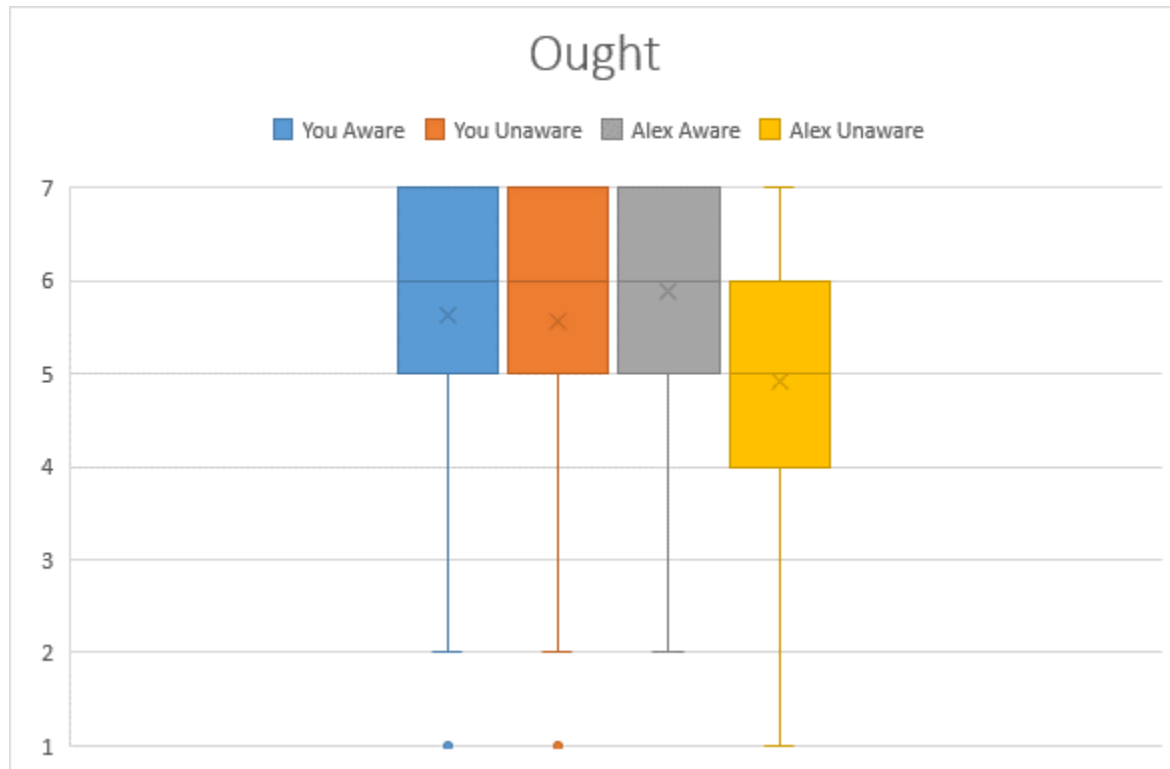
The data we collected is not a complete, traditional Likert scale. It is a collection of Likert-type items with values ranging over 1 (“Strongly Disagree”) to 7 (“Strongly Agree”).¹⁶ Because it is Likert-type, and not full Likert scale data, the items cannot be assumed to have a normal distribution or a ratio data type.¹⁷ Likert-type items are of an ordinal type; this means the values have an order, but the distances between the values are not necessarily equal to each other. Ordinal data limits the types of tests we can use because common statistics, such as mean and variance, have no meaning for ordinal data. Meaningful statistics for ordinal data include medians, modes, frequencies, and quartiles.

With that in mind, we first looked at attributions of moral obligation (ought). Overall, participants agree that the agent (either “You” or “Alex”) ought to give up the seat independent of the agent’s situational awareness. The most common judgment for vignettes in which the agent (either “You” or “Alex”) is situationally unaware is 6 for “agree,” whereas the most common judgment for vignettes in which the agent (either “You” or “Alex”) is able is 7 for “strongly agree” (see Figure 1).

¹⁶ We chose not to dichotomize the data into two categories--agree/disagree, or for example blame/don’t blame--because we want to be able to test the strength of the agreement. Dichotomizing would mean we would not be able to discover if one receives strong blame and another is only somewhat blameworthy. We would only be able to say that blame was assigned in both cases. For more on the problems with dichotomization, see MacCallum et al (2002).

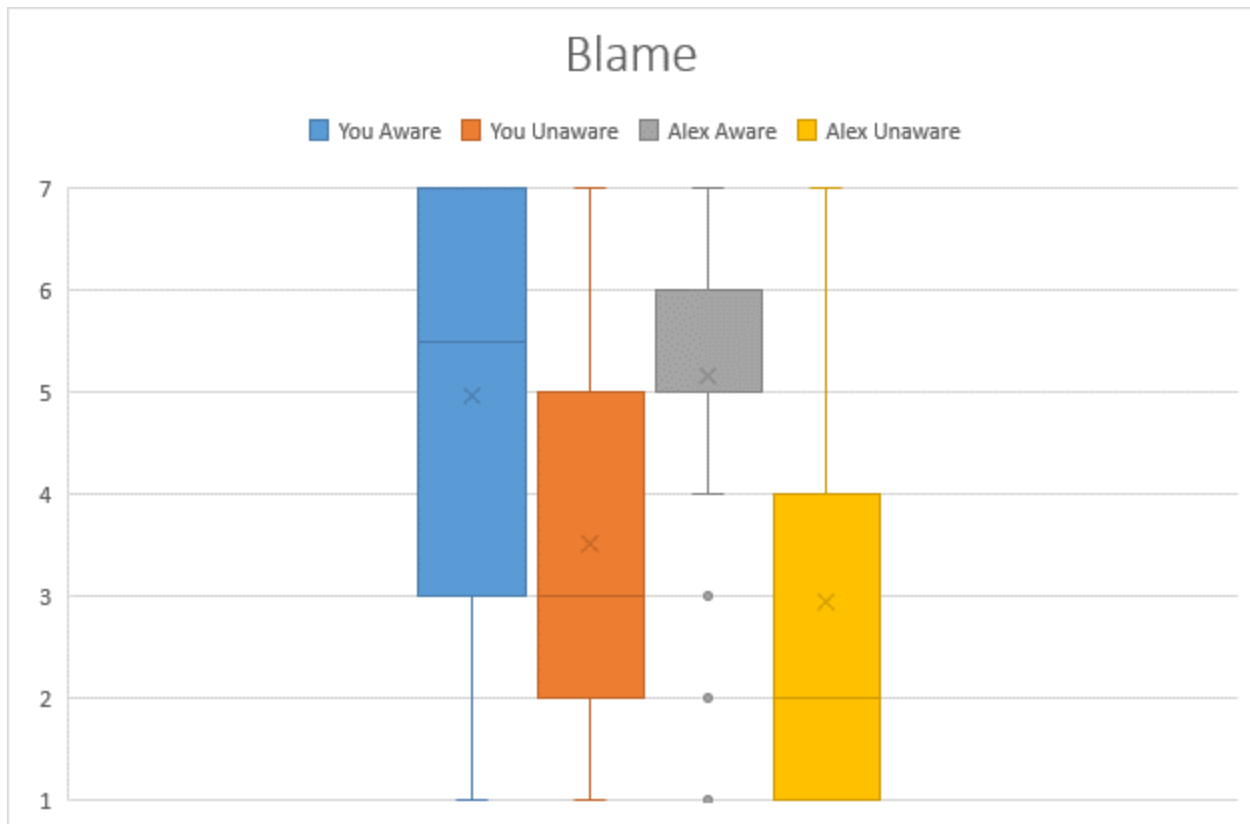
¹⁷ Likert-type items are single questions that still use the same layout as the questions within a full Likert scale. Individual items in a full Likert-scale are then summed to form what is assumed to be a continuous variable. For more on the statistical difference between Likert-type items and full Likert scales, see Clason and Dormody (1994).

Figure 1. Distribution of agreement with “ought” statements.



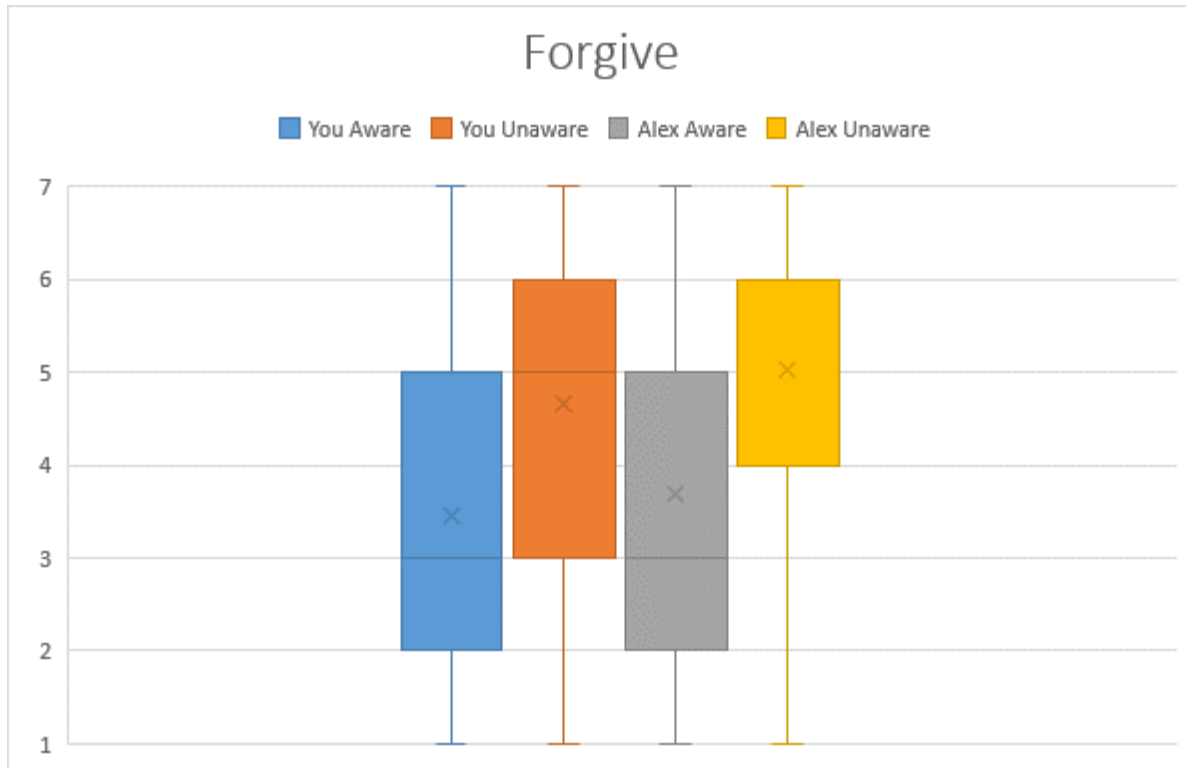
Unlike attributions of moral obligation, attributions of blame were more dependent on circumstances (see Figure 2). When the agent (either “You” or “Alex”) is situationally aware, the most common responses are ones of agreement with attributing blame. On the other hand, when the agent is situationally unaware, the most common responses were to disagree with attributing blame.

Figure 2. Distribution of agreement with “blame” statements.



Attributions of forgiveness followed surprisingly similar trends as attributions of blame (see Figure 3). Usually, vignettes including situationally aware agents garnered less forgiveness than those with situationally unaware agents. The most common responses for situationally aware agents are disagreement with forgiveness. For situationally unaware agents, respondents most commonly agreed with forgiving the agent.

Figure 3. Distribution of agreement with “forgive” statements.



In addition to analysis of spread, which indicates overall agreement with statements of moral obligation, blame, and forgiveness, we analyzed the main effects and interactions effects in our 2x2 factorial design. Again, because of limitations on Likert-type ordinal data, we used a less common and more specialized type of analysis. We chose to perform an Aligned Rank Transform¹⁸ (ART) on our response data. ART produces ratio statistics from ordinal data that can then be analyzed using traditional ANOVA tests. This method allowed us to identify the significant effects in each of the four statements.

We checked for two main effects and their interaction effect. We found that interaction did not have a significant effect in any of the cases. Attributions of moral obligation were affected by situational awareness ($p = 6.19 \times 10^{-05}$) and agent ($p = 0.017$). Attributions of blame

¹⁸ It is common to test Likert-type data as if it is continuous. The commonality of this practice does not mean it is the appropriate test. For example, Likert-type data fails the assumptions for ANOVA testing meaning the results from such a test may be wrong. For a detailed explanation of ART procedures, see Wobbrock et al. (2011). We performed the statistical analysis in R using the ARTool package provided by Wobbrock et al. An overview of the tool and package can be accessed at <https://depts.washington.edu/aimgroup/proj/art/>.

were not affected by agent but they were affected by situational awareness ($p = 2.35 \times 10^{-13}$). Attributions of forgiveness were also affected by situational awareness ($p = 3.34 \times 10^{-09}$) but not by agent. As we mentioned in Section 2a, demographic factors had no significant effects on participants' responses.

Since there were no interaction effects, we investigated separate groups according to main effects only and compared these groups. Keeping in mind the ordinal nature of our data, we chose to perform Wilcoxon Rank Sum tests.¹⁹ Only situational awareness, and not agent, affected “blame” and “forgive” judgments. We found that participants blame themselves ($p = 1.48 \times 10^{-05}$) and Alex ($p = 8.22 \times 10^{-09}$) more in vignettes where the agent is situationally aware. Participants forgive themselves more in vignettes where they are situationally unaware ($p = 4.47 \times 10^{-05}$) and forgive “Alex” more in vignettes where “Alex” is situationally unaware ($p = 1.55 \times 10^{-05}$). Overall, when comparing “blame” and “forgive” judgments, participants assign more blame than forgiveness to themselves when they are situationally aware ($p = 6.15 \times 10^{-07}$) and more forgiveness than blame to themselves when they are situationally unaware ($p = 3.80 \times 10^{-04}$). Similar judgments were given when the agent is “Alex,” with more blame than forgiveness in vignettes where the agent is situationally aware ($p = 9.95 \times 10^{-06}$) and more forgiveness than blame in vignettes where the agent is situationally unaware ($p = 3.10 \times 10^{-08}$). These comparisons indicate that in situations where the agent is situationally unaware, the agent is likely to be forgiven for whatever blame the agent was originally assigned. On the other hand, when the agent is situationally aware of the facts surrounding the (in)action for which the agent is blamed, the agent is unlikely to be completely forgiven.

In addition to the data set we have analyzed so far, we also looked at data from pilot studies. These pilot studies were originally given to determine the best wording for the vignettes

¹⁹ The null hypothesis for these tests is that “true location shift is equal to 0” for two between-subjects groups.

used in the main study. These pilot studies only focused on vignettes where the agent is “You.” There were a very small number of responses to these studies ($n = 64$). We mention them here because, despite the small sample size and limitations, the results agree with those of the larger survey. Situational awareness did not affect participants’ judgments of obligation ($p = 0.125$). Participants blamed themselves more when situationally aware ($p = 0.023$) and forgave themselves more when situationally unaware ($p = 0.004$).

To sum up, then, the main results of our experimental study are as follows:

1. When the moral agent (either “You” or “Alex”) is situationally aware, it is more likely that the agent will be blamed for failing to perform the obligatory action than when the agent is situationally unaware.
2. When the moral agent (either “You” or “Alex”) is situationally aware, it is less likely that the agent will be forgiven for failing to perform the obligatory action than when the agent is situationally unaware.
3. When the moral agent is oneself (“You”) and one (“You”) is situationally aware, it is more likely that the agent (“You”) be blamed than forgiven.
4. When the moral agent is oneself (“You”) and one (“You”) is situationally unaware, people forgive the agent (“You”) more than they blame the agent (“You”).
5. When the moral agent is not oneself (“Alex”) and the agent (“Alex”) is situationally aware, people blame the agent (“Alex”) more than they forgive the agent (“Alex”).
6. When the moral agent is not oneself (“Alex”) and the agent (“Alex”) is situationally unaware, people forgive the agent (“Alex”) more than they blame the agent (“Alex”).

Accordingly, our results suggest that whether a moral agent (the subject that can be worthy of blame or forgiveness) is situationally aware or unaware does have an effect on people’s

attributions of blame and forgiveness (regardless of consequences or side effects). In general, it was more likely that a situationally aware agent will be blamed for failing to perform the obligatory action than a situationally unaware agent. Furthermore, in general, it was less likely that a situationally aware agent will be forgiven for failing to perform the obligatory action than a situationally unaware agent. When the agent is situationally unaware, it is more likely that the agent will be forgiven than blamed.

2c. Discussion

As we mentioned in Section 1, many philosophers find it intuitive that moral responsibility has an epistemic dimension. More explicitly, whether a moral agent is worthy of praise or blame is thought to depend, at least in part, on the moral agent's epistemic situation. For instance, Fischer and Ravizza (1998, p. 13) claim "that an agent is responsible only if he both knows the particular facts surrounding his action, and acts with the proper sorts of beliefs and intentions." Fischer and Ravizza (1998, p. 12) consider the case of "the person who backs his car out of his garage unaware that a tiny kitten is snoozing beneath the rear tire," and say that "it would be odd to judge that the driver is morally responsible for the kitten's untimely death." As Mason (2015, p. 3038) puts it:

it seems as though ignorance should exculpate, as it can in non-moral cases. *Non-culpable ignorance of fact is usually a straightforward excuse*. If I didn't know that the lever was connected to a puppy killing device, and there is no way I could have known or suspected that it was, then I am not blameworthy for pulling the lever (emphasis added).

Likewise, according to Harman (2011, p. 444), it is "natural [to think] that ignorance of non-moral matters can exculpate." Now, we think that our results provide some empirical support for the hypothesis that there is something intuitive or "natural" about the "thought that ignorance of

non-moral matters can exculpate” (Harman 2011, p. 444). For, in our vignettes, the moral agent is an *ignorant* agent,²⁰ since the agent is asleep, and so is unaware of the facts surrounding his/her (in)action, specifically, that a person using crutches just boarded the train and there are no available seats. The fact that it is less likely that the ignorant agent will be blamed suggests that they take the agent’s situational unawareness (i.e., that the agent is unaware of the facts that a person using crutches just boarded the train and that there are no available seats) as a reason to withhold blame. It is important to emphasize that participants also judge that the agent *ought* to give up the seat for the person using crutches. This means that participants think that the agent has failed to do what is morally obligatory, and yet they are still willing to withhold blame when the agent is unaware of the facts that a person using crutches just boarded the train and that there are no available seats. When the moral agent is situationally aware, i.e., when the agent is aware of the facts that a person using crutches just boarded the train and that there are no available seats, it is more likely that the situationally aware agent will be blamed for failing to give up the seat.

Moreover, our results suggest that what holds for attributions of blame probably holds for attributions of forgiveness as well. For, as we argue above, the moral agent in our vignettes is an *ignorant* agent,²¹ since the agent is asleep, and so is unaware of the facts surrounding his/her (in)action, specifically, that a person using crutches just boarded the train and there are no available seats. The fact that it is more likely that the ignorant agent will be forgiven suggests that they take the agent’s situational unawareness (i.e., that the agent is unaware of the facts that a person using crutches just boarded the train and that there are no available seats) as a reason to

²⁰ It is important to note here that the agent is not morally ignorant, i.e., the agent is not ignorant of moral facts. Rather, the agent is situationally ignorant, i.e., ignorant of the non-moral facts surrounding his/her (in)action, in particular, that that a person using crutches just boarded the train and there are no available seats. On moral ignorance, see Guerrero (2007).

²¹ See footnotes 13 and 20.

forgive the agent. Again, it is important to emphasize that participants also judge that the agent *ought* to give up the seat for the person using crutches. This means that participants think that the agent has failed to do what is morally obligatory, and yet they are still willing to forgive the agent when the agent is unaware of the facts that a person using crutches just boarded the train and that there are no available seats. When the moral agent is situationally aware, i.e., when the agent is aware of the facts that a person using crutches just boarded the train and that there are no available seats, it is less likely that the situationally aware agent will be forgiven for failing to give up the seat.

These results lend some empirical support to the hypothesis that it is “intuitive” or “natural” to think that non-moral or factual ignorance (specifically, situational unawareness) exculpates for the following reasons. As we have seen, our results suggest that, when the agent is situationally aware, people are more inclined to blame and less inclined to forgive the agent for failing to perform the obligatory action, but when the agent is situationally unaware, people are more inclined to forgive than blame the agent for failing to perform the obligatory action. In other words, in the vignettes in which the agent is situationally aware, it *seems* to people that the agent is worthy of blame but not forgiveness, whereas in the vignettes in which the agent is situationally unaware, it *seems* to people that the agent is more worthy of forgiveness than blame. Given that intuitions are intellectual seemings (Brogaard 2014), it follows that people find it *intuitive* that situationally aware agents are worthy of blame but not forgiveness, whereas situationally unaware agents are more worthy of forgiveness than blame. In other words, our results suggest that people find the thought that non-moral or factual ignorance can exculpate *intuitive*, just as many professional philosophers do (see, e.g., Harman 2011 and Mason 2015),

given that they intuitively judge that situationally aware agents are more worthy of blame than forgiveness, whereas situationally unaware agents are more worthy of forgiveness than blame.

3. Objections

We have argued that our results lend some empirical support to the hypothesis that it is “intuitive” or “natural” to think that non-moral or factual ignorance (specifically, situational unawareness) exculpates. In this section, we discuss a few objections to this interpretation of our results.

An anonymous reviewer keenly observed that, given our between-subjects experimental design, it is possible that some participants were either reluctant to blame, or willing to forgive, the situationally unaware agent for not giving up the seat because they thought that the agent is under no moral obligation to do in the first place or because the agent lacks the ability to do so. In other words, according to the reviewer, our results are consistent with two alternative hypotheses: (a) participants were either reluctant to blame, or willing to forgive, the situationally unaware agent for not giving up the seat because they did not think that the agent ought to give up the seat; and (b) participants were either reluctant to blame, or willing to forgive, the situationally unaware agent for not giving up the seat because they did not think that the agent is able to give up the seat.

In reply, we would like to make the following points. First, our data on “ought” statements show that participants generally think that the agent ought to give up the seat. Even the lowest of the “ought” judgments for Alex Unaware is “somewhat agree” (see Figure 1). These experimental results count as some empirical evidence against alternative hypothesis (a),

which would predict—contrary to our results—that most participants would *not* agree with “ought” statements.

Second, recent experimental studies on the principle known as “*Ought Implies Can*,” which states that an agent ought to do something only if she can do it, suggest that participants are generally willing to assign moral obligation to agents even when those agents are unable to perform the obligatory action in question.²² These experimental results, then, count as some empirical evidence against alternative hypothesis (b), which would predict—contrary to recent results of OIC studies—that most participants would not assign moral obligation to unable agents.

Third, even if it is possible that participants were either reluctant to blame, or willing to forgive, the situationally unaware agent for not giving up the seat because they thought that the agent is under no moral obligation to do in the first place, it is still the case that there is an *epistemic* dimension here. That is, it is a non-moral, epistemic fact about the situational awareness (or lack of awareness) of an agent that seems to have an effect of attributions of blame and forgiveness. In other words, even if alternative hypothesis (b) were true, it would still be the case that situational awareness (or unawareness) is that which has an effect on our participants’ attributions of blame or forgiveness to the agent (You or Alex), since the only reason to think that the agent is unable to give up the seat is the non-moral facts that the agent is sleeping. Our data do not support any claims about a particular relation between moral obligation and blame (or moral obligation and forgiveness). But they do support the claim that an epistemic dimension, specifically, facts about an agent’s situational awareness (or lack of awareness), play a role in mitigating blame and increasing forgiveness. Further studies are needed to understand the role of non-moral epistemic facts in mitigating blame and increasing forgiveness in addition to the

²² See Buckwalter and Turri (2015), Mizrahi (2015a), (2015b), Chituc et al. (2016), and Henne et al. (2016).

substantial philosophical issue of whether or not obligations persist after assignments of blame and/or forgiveness.

It might seem as if these results are expected, and thus not surprising at all. In one sense, our results concerning blame attributions are to be expected, given that most philosophers do find it intuitive that non-moral, factual ignorance can exculpate, as we discussed in Section 1. In another sense, however, our results are unexpected if considered relative to recent results in experimental philosophy. Many studies in experimental philosophy so far have reported “negative” results; that is, their results suggest that non-philosophers generally do not share the intuitions that professional philosophers typically have. This has led many critics of experimental philosophy to complain that experimental philosophy is mostly “negative” (see, e.g., Talbot 2014). Indeed, even experimental philosophers talk about “the negative program” in experimental philosophy; that is, that (at least some) work in experimental philosophy seeks to undermine the evidential value of intuitions (Sytsma and Livengood 2016, p. 51). Accordingly, to find out that the intuitive judgments of professional philosophers and non-philosophers sometimes align is somewhat surprising, in light of the mostly “negative” findings that came out of experimental philosophy so far. In other words, experimental results in philosophy, which typically show a divergence of opinions between professional philosophers and non-philosophers, now show convergence with respect to whether situational unawareness exculpates.²³

Moreover, such “positive” results also contribute in a meaningful way. Both moral psychologists and experimental philosophers have to contend with problems related to the idiosyncrasies of research involving human subjects and statistical methodologies. One of the

²³ For more on this sort of “positive” program, or more precisely, experimental philosophy as cognitive science, see Knobe (2016).

most recent and prominent examples of this is the so-called “replication crisis” (see, e.g., Open Science Collaboration 2015). For this reason, the fact that our results in both large and small sample studies, which are derived from somewhat unpopular but specialized statistical analyses, corroborate other results is no small matter.

Second, the results about forgiveness attributions are somewhat unexpected insofar as they are “positive,” like the results about blame attributions, and they suggest that there may be a symmetry between being worthy of blame, or blameworthiness, and being worthy of forgiveness. When philosophers talk about moral responsibility, they usually talk about praise and blame and/or reward and punishment (Fischer 2006, p. 63).²⁴ This is not to say that philosophers have nothing to say about forgiveness (see, e.g., Allais 2008 and Holmgren 2012), even about forgiveness and blame (see, e.g., Kekes 2009 and Gamlund 2011). Nevertheless, our results suggest that perhaps experimental philosophers and moral psychologists who are interested in moral responsibility should give more *empirical* attention to forgiveness.²⁵ In particular, we propose that our results may prove theoretically fruitful in addressing a problem that Warmke and McKenna (2013, p. 195) identify, namely, “what philosophical strategies should be employed for determining which view [of forgiveness] is preferable?” If we forgive others in different contexts, in response to different kinds of wrongs, and for different reasons, as an anonymous reviewer helpfully pointed out, then an experimental approach can be useful in terms of identifying these various contexts, or so we think.²⁶ The task of understanding forgiveness has

²⁴ Strawson (1962) does include forgiveness in his list of “reactive attitudes,” which also includes resentment, indignation, anger, and gratitude. As mentioned above, however, philosophers of moral responsibility have focused mostly on praise and blame. See, e.g., Macnamara (2011).

²⁵ See Martin and Cushman (2016) for an experimental test of the hypothesis that we forgive unintentional actions because “we do not believe the agent is even causally responsible for them” (Martin and Cushman 2016, p. 134). Recall that our experimental study is not about intentional (or unintentional) actions but rather about situational awareness (or lack of situational awareness) and attributions of blame as well as forgiveness.

²⁶ As another anonymous reviewer helpfully pointed out, a possible symmetry between blameworthiness and being worthy of forgiveness might be taken as a platitude that any account of forgiveness must be able to accommodate.

become especially urgent since the end of the twentieth century, according to Moody-Adams (2015, p, 161), “as divided societies looked to forgiveness as a vehicle of reconciliation, governments sought forgiveness for past wrongs, and popular psychology explored the therapeutic effects of forgiveness.”

As an anonymous reviewer pointed out, contrary to the participants in our study, it might strike some as a bit odd to say that the sleeping agent is worthy of blame. As stipulated in our vignettes, the sleeping agent does not give up his/her seat to the person on crutches despite the fact that there are no available seats on the train. However, it might seem to some, it is not true that the sleeping agent omitted giving up his/her seat, which means that the sleeping agent is not blameworthy of not giving up his/her seat.

We would like to make two points in reply. First, the reviewer’s and our participants’ intuition seem to be in agreement here. For, when the agent (either “You” or “Alex”) is situationally unaware (i.e., sleeping), our participants agree with forgiving the agent more than they agree with blaming the agent.

Second, the anonymous reviewer’s objection raises an important question about the difference, if any, between “not doing” and “omitting.” It is clear that our sleeping agent did not give up the seat to the person on crutches. But does that mean that the sleeping agent omitted giving up the seat? This is not an easy question to answer, and doing justice to it is clearly beyond the scope of this paper. On the one hand, it might be thought that omission implies trying to do something but failing. After all, *The Oxford English Dictionary* definition of “omission” is “a failure to fulfill a moral or legal obligation.” If we take “failure” here to mean, “trying but failing,” then we could say that “omitting to A” means “trying to A but failing to do so.” On the

However, we think that it would be useful to have empirical evidence, even for something that is considered a platitude, when possible.

other hand, we also talk about failing to do something even without trying. For example, it is natural to say things like “he failed to follow her advice” even if he did not even try to follow her advice. Similarly, it is natural to say, “Michael Flynn failed to disclose a trip he took to the Middle East to explore a business deal with the Saudi government and a Russian government agency” (CNN Politics 2017), even if he did not even try to disclose that information. Based on somewhat similar considerations, Clarke (2012, p. 140) concludes that “in very many cases an omission [...] is simply an absence of action.” According to Clarke (2012, p. 140), these are the cases in which the absent action “was called for by some norm or standard.” In our vignettes, there is clearly an absence of action on the part of the sleeping agent, i.e., the sleeping agent did not give up the seat to the person on crutches, which is called for by some norm. The norm is that of giving up seats on public transportation to people with disabilities and/or injuries. For most of our participants, this seemed sufficient to make the situationally unaware (or ignorant) agent more worthy of forgiveness than blame.

To this it might be objected, as an anonymous reviewer did, that the following are equally plausible alternatives to saying that it is situational awareness (or unawareness) that affects our participants’ attributions of blame or forgiveness to the agent (You or Alex):

- (i) Deny that the norm of giving up seats on public transportation to people with disabilities and/or injuries applies to sleeping people and argue that the agent is not morally obligated to give up the seat;
- (ii) Concede that the agent ought to give up the seat but deny that the agent omitted an obligatory action by remaining asleep;

- (iii) Argue that the agent cannot be blamed for not giving up the seat because the omission is unintentional, i.e., the agent did not choose to remain seated just as she did not choose to fall asleep.

We think that the reviewer is correct in pointing out that these are equally plausible alternatives. But we also think that our results do provide some empirical evidence against them. On (i), we would expect most of our participants to disagree with “ought” statements. As we have seen, however, most of our participants agree with “ought” statements. Even the lowest of the “ought” judgments for Alex Unaware is “somewhat agree” (see Figure 1). In other words, most of our participants agree that the agent has not done what he or she ought to have done, namely, to give up the seat on a train for a disabled person. Moreover, even if the norm of giving up seats on public transportation to people with disabilities and/or injuries had an exclusion clause along the lines of “Sleeping passengers need not get up for people with disabilities,” the exception would still be based on situational awareness, or more precisely unawareness, for it excludes *sleeping* passengers. Why else would sleeping passengers be excluded from this norm if not for the non-moral fact that they are unaware of relevant non-moral facts (e.g., that there is a disabled person that needs a seat)? In other words, it would still be situational unawareness that seems to be doing the exculpation in this case.

On (ii), the agent ought to give up the seat, but the fact that she did not give up the seat does not count as an omission of an obligatory action. But why? If it is conceded that all passengers, including sleeping ones, ought to give up their seats to people with disabilities, why does it matter whether we say that one *failed* to do or *omitted* to do what one ought to? What matters is whether the seat had been given up, as it should have. In our vignettes, it is clear that the seat had not been vacated for the person with disabilities to use. Moreover, our data show

that most of our participants agree that the agent ought to give up the seat. Therefore, as most of our participants seem to see it, the agent did not do what she ought to have done.

On (iii), the idea seems to be that the agent is not culpable for falling asleep, and hence she is also not culpable for failing to give up the seat. As far as our vignettes are concerned, however, we can see how participants might think of the agent's falling asleep (and, consequently, lacking in situational awareness) as either culpable or non-culpable. For example, one who thinks of the agent as riding the train in the evening, going back home after a long day at work, and finally falling asleep despite efforts to stay awake, might think of the agent's ignorance of the non-moral facts that a person using crutches just boarded the train and there are no available seats as non-culpable. On the other hand, one who thinks of the agent as riding the train in the morning, fresh after a good night's sleep, but still falls asleep while listening to music while trying to avoid eye contact with other passengers, might think of the agent's ignorance of the non-moral facts that a person using crutches just boarded the train and there are no available seats as culpable. Our study does not control for this, which means that further studies are needed to find out how culpability (or lack thereof) might affect attributions of blame and forgiveness in similar cases. In that respect, the reviewer's comment is very helpful in terms of suggesting ways to test that in future studies.

4. Conclusion

In this paper, we set out to test empirically an idea that many philosophers find intuitive, namely, that moral responsibility has an epistemic dimension having to do with the moral agent's epistemic situation. Our results show that whether a moral agent (the subject of blame or forgiveness) is aware of the facts surrounding his/her inaction (e.g., failing to give up a seat on a

train for a person using crutches) does have an effect on people's attributions of blame. In general, it was more likely that a situationally aware agent will be blamed for failing to perform the obligatory action than a situationally unaware agent. We also tested attributions of forgiveness in addition to attributions of blame. In general, it was less likely that a situationally aware agent will be forgiven for failing to perform the obligatory action than a situationally unaware agent. When the agent is situationally unaware, it is more likely that the agent will be forgiven than blamed.

These results suggest that people find the thought that non-moral or factual ignorance can exculpate *intuitive*, just as many professional philosophers do (see, e.g., Harman 2011 and Mason 2015), given that they intuitively judge that situationally aware agents are more worthy of blame than forgiveness, whereas situationally unaware agents are more worthy of forgiveness than blame. For this reason, we have argued, these results lend some empirical support to the hypothesis that there is something intuitive or "natural" about the "thought that ignorance of non-moral matters can exculpate" (Harman 2011, p. 444). We have also proposed that further experimental studies are needed to explore a possible symmetry between blameworthiness and forgiveness-worthiness, which is suggested by our findings, and that may be theoretically fruitful as far as the ordinary concept of moral responsibility is concerned.

Acknowledgments

We are grateful to two anonymous reviewers of *Acta Analytica* for helpful comments on earlier drafts of this paper.

References

Agule, C. K. (2016). Resisting Tracing's Siren Song. *Journal of Ethics & Social Philosophy* 10 (1): 1-24.

Allais, L. (2008). Wiping the Slate Clean: The Heart of Forgiveness. *Philosophy and Public Affairs* 36 (1): 33-68.

Baz, A. (2016). Recent Attempts to Defend the Philosophical Method of Cases and the Linguistic Turn. *Philosophy and Phenomenological Research* 92 (1): 105-130.

Benchimol, J. (2011). The Moral Significance of Unintentional Omission: Comparing Will-Centered and Non-will-centered Accounts of Moral Responsibility. In N. A. Vincent, I. van de Poel, and J. van den Hoven (eds.), *Moral Responsibility: Beyond Free Will and Determinism*. (pp. 101-120). Dordrecht: Springer.

Brekke Carlsson, A. (2017). Blameworthiness as Deserved Guilt. *The Journal of Ethics* 21 (1): 89-115.

Brogaard, B. (2014). Intuitions as Intellectual Seemings. *Analytic Philosophy* 55 (4): 382-393.

Buckwalter, W. and Turri, J. (2015). Inability and Obligation in Moral Judgment. *PLoS ONE* 10 (8): 1-20.

Chituc, V., Henne, P., Sinnott-Armstrong, W., and De Brigard, F. (2016). Blame, not ability, impacts moral “ought” judgments for impossible actions: Toward an empirical refutation of “ought” implies “can”. *Cognition* 150:20-25.

Clarke, R. (2012). What is an Omission? *Philosophical Issues* 22 (1): 127-143.

Clason, D. and Dormody, T. (1994). Analyzing Data Measured by Individual Likert-Type Items. *Journal of Agricultural Education* 35 (4): 31-35.

Cole, J. and Bengson, J. (2009). Asymmetries in Judgments of Responsibility and Intentional Action. *Mind & Language* 24 (1): 24-50.

Cova, F. (2016). The Folk Concept of Intentional Action: Empirical Approaches. In J. Sytsma and W. Buckwalter (eds.), *A Companion to Experimental Philosophy* (pp. 121-141). Malden, MA: Wiley Blackwell.

Cova, F. (2017). Frankfurt-style cases and the explanation condition for moral responsibility: a reply to Swenson. *Acta Analytica*. doi:10.1007/s12136-017-0316-5.

Cushman, F. (2008). Crime and Punishment: Distinguishing the Roles of Causal and Intentional Analyses in Moral Judgment. *Cognition* 108 (2): 353-380.

Faraci, D. and Shoemaker, D. (2014). Huck vs. Jojo: Moral Ignorance and the (A)symmetry of Praise and Blame. In J. Knobe, T. Lombrozo, and S. Nichols (eds.), *Oxford Studies in Experimental Philosophy* (pp. 7-27). New York: Oxford University Press.

Feltz, A. (2007). The Knobe Effect: A Brief Overview. *Journal of Mind and Behavior* 28 (3-4): 265-277.

Fischer, J. (2003). 'Ought-implies-can', causal determinism and moral responsibility. *Analysis* 63 (3): 244-250.

Fischer, J. (2006). *My Way: Essays on Moral Responsibility*. New York: Oxford University Press.

Fischer, J., and Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility*. New York: Cambridge University Press.

Fischer, J., and Tognazzini, N. (2009). The Truth about Tracing. *Nous* 43 (3): 531-556.

Gamlund, E. (2011). Forgiveness without Blame. In C. Fricke (ed.), *The Ethics of Forgiveness: A Collection of Essays* (pp. 107-129). New York: Routledge.

Ginet, C. (2000). The Epistemic Requirements for Moral Responsibility. *Nous* 34 (14): 267-277.

Graham, P. A. (2011). Fischer on Blameworthiness and “Ought” Implies “Can.” *Social Theory and Practice* 37 (1): 63-80.

Guerrero, A. (2007). Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution. *Philosophical Studies* 136 (1): 59-97.

Haji, I. (1997). An Epistemic Dimension of Blameworthiness. *Philosophy and Phenomenological Research* 57 (3): 523-544.

Harman, E. (2011). Does Moral Ignorance Exculpate? *Ratio* 24 (4): 443-468.

Henne, P., Chituc, V., De Brigard, F., and Sinnott-Armstrong, W. (2016). An Empirical Refutation of ‘Ought’ Implies ‘Can’. *Analysis* 76 (3): 283-290.

Holmgren, M. R. (2012). *Forgiveness and Retribution: Responding to Wrongdoing*. New York: Cambridge University Press.

Kekes, J. (2009). Blame versus Forgiveness. *The Monist* 92 (4): 488-506.

Khoury, A. C. (2012). Responsibility, Tracing, and Consequences. *Canadian Journal of Philosophy* 42 (3): 187-208.

- Knapp, T. (1990). Treating Ordinal Scales as Interval Scales: An Attempt to Resolve the Controversy. *Nursing Research* 39 (2): 121-123.
- Knobe, J. (2003a). Intentional Action in Folk Psychology: An Experimental Investigation. *Philosophical Psychology* 16 (2): 309-325.
- Knobe, J. (2003b). Intentional Action and Side Effects in Ordinary Language. *Analysis* 63 (3): 190-194.
- Knobe, J. (2004a). Intention, Intentional Action and Moral Considerations. *Analysis* 64 (2): 181-187.
- Knobe, J. (2004b). Folk Psychology and Folk Morality: Response to Critics. *Journal of Theoretical and Philosophical Psychology* 24 (2): 270-279.
- Knobe, J. (2016). Experimental Philosophy is Cognitive Science. In J. Sytma & W. Buckwalter (Eds.), *A Companion to Experimental Philosophy* (pp. 37-52). Malden, MA: Wiley Blackwell.
- Lagnado, D. A. and Channon, S. (2008). Judgments of Cause and Blame: The Effects of Intentionality and Foreseeability. *Cognition* 108 (3): 754-770.
- MacCallum, R. C., Zhang, S., Preacher, K. J., and Rucker, D. (2002). On the Practice of Dichotomization of Quantitative Variables. *Psychological Methods* 7 (1): 19-40.

Macnamara, C. (2011). Holding Others Responsible. *Philosophical Studies* 152 (1): 81-102.

Martin, J. W. and Cushman, F. (2016). Why We Forgive What Can't Be Controlled? *Cognition* 147: 133-143.

Mason, E. (2015). Moral Ignorance and Blameworthiness. *Philosophical Studies* 172 (11): 3037-3057.

Mele, A. (2011). Moral Responsibility for Actions: Epistemic and Freedom Conditions. *Philosophical Explorations* 13 (2): 101-111.

Mizrahi, M. (2009). 'Ought' Does Not Imply 'Can'. *Philosophical Frontiers* 4 (1): 19-35.

Mizrahi, M. (2012). Does 'Ought' Imply 'Can' from an Epistemic Point of View? *Philosophia* 40 (4): 829-840.

Mizrahi, M. (2014). Does the Method of Cases Rest on a Mistake? *Review of Philosophy and Psychology* 5 (2): 183-197.

Mizrahi, M. (2015a). Ought, Can, and Presupposition: An Experimental Study. *Method: Analytic Perspectives* 4 (6): 232-243.

Mizrahi, M. (2015b). Ought, Can, and Presupposition: A Reply to Kurthy and Lawford-Smith. *Method: Analytic Perspectives* 4 (6): 250-256.

Montminy, M. (2016). Doing One's Reasonable Best: What Moral Responsibility Requires. *Journal of the American Philosophical Association* 2 (1): 55-73.

Moody-Adams, M. (2015). The Enigma of Forgiveness. *Journal of Value Inquiry* 49 (1): 161-180.

Nadelhoffer, T. (2004). On Praise, Side-Effects, and Folk Ascriptions of Intentionality. *Journal of Theoretical and Philosophical Psychology* 24 (2): 196-213.

Nagel, J. (2012). Intuitions and Experiments: A Defense of the Case Method in Epistemology. *Philosophy and Phenomenological Research* 85 (3): 495-527.

Open Science Collaboration. (2015). Estimating the Reproducibility of Psychological Science. *Science* 349(6251). DOI: 10.1126/science.aac4716.

Oshana, M. (2013). Second Hand (Moral) Responsibility in Law. In I. Haji and J. Caouette (eds.), *Free Will and Moral Responsibility* (pp. 229-248). Newcastle: Cambridge Scholars Publishing.

Rosen, G. (2004). Scepticism about Moral Responsibility. *Philosophical Perspectives* 18 (1): 295-313.

Rosen, G. (2008). Kleinbart the Oblivious and other Tales of Ignorance and Responsibility. *The Journal of Philosophy* 105 (10): 591-610.

Sauer, H. (2014). It's the Knobe Effect, Stupid! *Review of Philosophy and Psychology* 5 (4): 485-503.

Sher, G. (2009). *Who Knew? Responsibility without Awareness*. New York: Oxford University Press.

Strawson, P. F. (1962). Freedom and Resentment. *Proceedings of the British Academy* 48: 1-25.

Sytsma, J. and Livengood, J. (2016). *The Theory and Practice of Experimental Philosophy*. Tonawanda, NY: Broadview Press.

Talbot, B. (2014). Why So Negative? Evidence Aggregation and Armchair Philosophy. *Synthese* 191 (16): 3865-3896.

Tognazzini, N. A. (2010). Review of George Sher, *Who Knew? Responsibility without Awareness*. *Notre Dame Philosophical Review* (January 3, 2010): <http://ndpr.nd.edu/news/24254-who-knew-responsibility-without-awareness/>.

Vargas, M. (2005). The Trouble with Tracing. *Midwest Studies in Philosophy* 29 (1): 269-291.

Warmke, B. (2011). Is Forgiveness the Deliberate Refusal to Punish? *Journal of Moral Philosophy* 8 (4): 613-620.

Warmke and McKenna (2013). Moral Responsibility, Forgiveness and Conversation. In I. Haji and J. Caouette (eds.), *Free Will and Moral Responsibility* (pp. 189-212). Newcastle: Cambridge Scholars Publishing.

Wieland, J. W. (2015). What's Special about Moral Ignorance? *Ratio* 30 (2): 149-164.

Wieland, J. W. (forthcoming). The Epistemic Condition. In P. Robichaud and J. W. Wieland (eds.), *Responsibility - The Epistemic Condition*. Oxford University Press.

Wobbrock, J.O., Findlater, L., Gergle, D. and Higgins, J.J. (2011). The Aligned Rank Transform for nonparametric factorial analyses using only ANOVA procedures. *Proceedings of the ACM Conference on Human Factors in Computing Systems* (pp. 143-146).

Zaibert, L. (2009). The Paradox of Forgiveness. *Journal of Moral Philosophy* 6 (3): 365-393.