Manipulating Moral Dumbfounding: Inhibiting the Identification of Reasons

Cillian McHugh[1][0000−0002−9701−3232], Marek McGann[2], Eric R. Igou[1][0000−0001−7744−9648], and Elaine L. Kinsella[1][0000−0003−4835−8581]

[1] University of Limerick, Limerick, Ireland
[2] Mary Immaculate College, South Circular road, Limerick, Ireland
cillian.mchugh@ul.ie

Abstract. Moral dumbfounding occurs when people defend a moral judgement even though they cannot provide a reason in support of this judgement. It manifests as an admission of not having reasons, or the use of unsupported declarations ("it's just wrong") or tautological reasons ("because it's incest") as justifications for a judgment. It typically occurs for harmless taboos. It is cited as evidence for intuitionist or dual-process theories of moral judgement over rationalist approaches, however the phenomenon remains poorly understood. We test a conflict in dual-processes explanation of moral dumbfounding, where moral dumbfounding is an example of conflict between a habitual response (making a judgement) and a response that results from deliberation (providing a reason for the judgement). One prediction of this explanation is that under specific manipulations (e.g., cognitive load), responses should vary in predictable ways. The dumbfounding paradigm involves three possible responses: (a) providing reasons for a judgement (System 2/deliberative); (b) accepting the counter-arguments and rating the behaviour as "not wrong" (System 1/habitual); (c) a dumbfounded response (System 1/habitual). Cognitive load manipulations have been shown to inhibit deliberative responding. We present 4 studies in which dumbfounded responding was investigated under cognitive load manipulations. We hypothesised that rates of providing reasons would be reduced under cognitive load. The identification of reasons was inhibited in Studies 1 and 3, but not in Studies 2 and 4. The results do not provide strong evidence that moral dumbfounding be explained as conflict in dual-processes. Future research should investigate this further, addressing methodological limitations identified.

Keywords: moral judgement · dual-processes · moral dumbfounding.

1 Moral Dumbfounding: A Brief Overview

Moral dumbfounding occurs when people defend a moral judgement even though they cannot provide a reason in support of this judgement [9, 8, 12]. It has long been cited as evidence for intuitionist or dual-process theories of moral judgement [1, 2, 7, 8] over more rationalist approaches [11], and its discovery [9] coincided

with, and arguably contributed to the recent rise in intuitionist theories of moral judgements [8].

Despite the influence of moral dumbfounding on the morality literature, the phenomenon is not well understood. Until recently [13, 12] empirical evidence testing the phenomenon was limited to a single study ($N = 30$), unpublished in peer-review form. McHugh et al. [12], building on the original work [9], identified two responses that may be taken as indicators of moral dumbfounding. Firstly, people may explicitly admit to not having reasons for their judgment. Secondly, people may use unsupported declarations ("it's just wrong") or tautological reasons (naming the behaviour "because it's incest") as justifications for a judgment. The current research uses the methods developed by McHugh et al. [12] to test one potential explanation of moral dumbfounding.

2   Moral Dumbfounding as Conflict

Drawing on dual-process explanations of moral judgement [7, 2], the studies presented here aim to test the hypothesis that moral dumbfounding can be explained as conflict in dual-processes [5]. Conflicts occur when a habitual response is different from a response that results from deliberation. In the moral dumbfounding paradigm, the habitual response is to rate the behaviour as wrong, however, deliberation reveals that there may not be any reason for this judgement. This conflict can be resolved by further deliberation and the identification of alternative reasons for a judgement, or the over-riding of the habitual response and the changing of a judgement. In this view, moral dumbfounding is the failure to resolve, or the acknowledging of this conflict.

One prediction of explaining moral dumbfounding as conflict and adopting a dual-process model of moral judgement, is, that under specific manipulations, moral judgements should vary in predictable ways. In the dumbfounding paradigm, there are generally three possible responses: (a) providing reasons for a judgement; (b) accepting the counter-arguments and rating the behaviour as "not wrong"; (c) a dumbfounded response. These responses can be positioned in terms of dual-processes, whereby option (a) requires deliberation (traditionally System 2), and options (b) and (c) are grounded largely in habitual responding (traditionally System 1) [2, 1, 7].

3   Dumbfounding and Cognitive Load

Above we suggest that identifying reasons for a judgement requires more deliberation than accepting the counter-arguments, (accepting the "logical" argument, and adopting the judgement that is grounded in reasons is System 1/habitual response). We also hypothesised that identifying reasons requires more deliberation than a dumbfounded response (the initial judgement of the behaviour was habitual). Manipulations of cognitive load are known to inhibit deliberative responding, [6], and, as such, cognitive load should inhibit option (a) person's ability to provide reasons for their judgement leading to an increase in option (b) or (c)

(or both). That is, cognitive load should inhibit the identification of reasons for a judgement, leading to an increase in dumbfounding or an increase in accepting the counter-arguments and revising the judgement made.

We conducted four studies in which we tested the hypothesised relationship between cognitive load and dumbfounded responding. We hypothesised that a cognitive load manipulation would inhibit the identification of reasons, leading to higher rates of dumbfounding or changing judgements (or both).

4 Study 1 - College Sample

The aim of Study 1 was to investigate if a cognitive load manipulation influenced participants' ability to justify their judgement.

4.1 Study 1: Method

Participants and Design Study 1 was a between subjects design. The dependent variable was response to the critical slide (see below). The independent variable was cognitive load with two levels: present and absent [4].

A total sample of 66 participants (55 female, 11 male; $M_{age} = 22.42$, min $= 18$, max $= 57$, $SD = 6.86$) took part. Participants were undergraduate students, postgraduate students, and alumni from Mary Immaculate College (MIC), and University of Limerick (UL). Participation was voluntary and participants were not reimbursed for their participation.

Procedure and Materials Data were collected using an online questionnaire. Data collection took place in a designated computer laboratory in MIC. The experimenter remained in the laboratory for the duration of the study. Participants were first presented with an information sheet and consent form.

Participants in the experimental condition were presented with an eight digit number/letter string and asked to memorise the sequence. After 30 seconds, the experiment progressed to the next slide. Participants had the option to click "ok" and progress to the next slide after 15 seconds.

Participants were then presented with the "Julie and Mark" (*Incest*) vignette [9]. Participants rated how right or wrong the behaviour of Julie and Mark was, and were given an opportunity to provide reasons for their judgement. Following this, participants were presented with a series of counter-arguments, which refuted commonly used justifications for rating the behaviour as "wrong".

Dumbfounding was measured using the critical slide [12]. This contained a statement defending the behaviour and a question as to how the behaviour could be wrong ("Julie and Mark's behaviour did not harm anyone, how can there be anything wrong with what they did?"). There were three possible answer options: (a) "There is nothing wrong"; (b) an admission of not having reasons ("It's wrong but I can't think of a reason"); and finally a judgement with accompanying justification (c) "It's wrong and I can provide a valid reason". The order of these response options was randomised. Participants who selected (c)

were prompted to type a reason. The selecting of option (b), the admission of not having reasons, was taken to be a dumbfounded response. Following the critical slide, participants in the experimental condition were required to reproduce the eight digit number-letter string sequence presented previously. Following this a post-discussion questionnaire in which participants rated their response to the scenario across various dimensions [9].

4.2   Study 1: Results

Threre was no difference in initial judgement depending on cognitive load, $t(63.9)$ = 1.256, $p = 0.214$. The responses to the critical slide for the experimental group ($N = 33$) and the control group ($N = 33$) are displayed in Fig. 1. A chi-squared test for independence revealed a significant association between experimental condition and response to the critical slide, $\chi^2(2,\ N = 66) = 7.53$, $p = .023$, $V = .27$: under cognitive load more participants (15) selected "There is nothing wrong" than in the control group (5), the observed power power was .69.

Fig. 1. Study 1: Responses to critical slide for the experimental group ($N = 33$) and the control group ($N = 33$)

4.3   Study 1: Discussion

The aim of the Study 1 was to investigate if dumbfounded responding was influenced by cognitive load. Specifically, adopting a dual-systems model of moral judgement, it was hypothesised that cognitive load would lead to reduced levels of deliberative responding, leading to a reduction in successfully identifying reasons

for judgements. This may lead to (a) increased levels of dumbfounding, or (b) increased selecting of the "nothing wrong" response. As predicted we found higher rates of selecting of the "nothing wrong" response in the cognitive load group compared to the control group.

5   Study 2 - Online Replication

Study 1 demonstrated interesting variability in responses to the critical slide depending on cognitive load. The aim of Study 2 was to assess the replicability of the results of Study 1, using an online sample. In Study 1, the experimenter was in the room with the participants. This made it more difficult for participants to cheat on the memory task. This is not possible with an online sample. An alternative cognitive load manipulation was taken from De Neys and Schaeken [3], whereby a dot pattern is briefly presented to participants, and participants are required to reproduce the dot pattern at a later stage.

5.1   Study 2: Method

Participants and Design Study 2 was a between subjects design. The dependent variable was response to the critical slide. The independent variable was cognitive load with two levels: high and low.

A total sample of 100 participant (56 female, 44 male; $M_{age} = 38.38$, min $= 19$, max $= 72$, $SD = 12.41$) took part. Participants in this sample were recruited using Amazon's MTurk. Participants were paid $0.50 for their participation. Participants were recruited from English speaking countries or from countries where residents generally have a high level of English (e.g., The Netherlands, Denmark, Sweden).

Procedure and Materials Data were collected using an online questionnaire. Materials were largely the same as in Study 1, with a change to the cognitive load manipulation. Cognitive load was manipulated using a dot-pattern memory task [3].

Participants were presented with a 3 x 3 grid containing a dot pattern. This image disappeared after one second. Participants then answered a question relating to the moral judgement task. Following this, participants were asked to reproduce the dot-pattern. All participants took part in the memory task, and cognitive load was manipulated by varying the complexity of the patterns presented [3]. The control group were presented with simple patterns, containing three dots in a line, while the experimental group were presented with more complex dot patterns containing 4 dots, see Fig. 2.

Study 2 proceeded in much the same way as Study 1. There were four target questions during which participants were engaged in the memory task. A different pattern was presented before each of the following: the initial judgement, the initial opportunity to provide reasons, the critical slide, and the revised judgement. After each of these questions participant were required to reproduce the pattern. As in Study 1, dumbfounding was measured using the critical slide.

Fig. 2. Sample dot patterns - more simple for the control group (a) and higher complexity for the experimental condition (b)

5.2   Study 2: Results

Threre was no difference in initial judgement depending on cognitive load, $t(97.77)$ = 0.962, $p = 0.339$. The responses to the critical slide for the experimental group ($N = 51$) and the control group ($N = 49$) group are displayed in Fig. 3. A chi-squared test for independence revealed no association between experimental condition and response to the critical slide, $\chi^2(2, N = 100) = 1.07$, $p = .585$, V = .10. The observed power was .14.

Fig. 3. Study 2: Responses to critical slide for the experimental group ($N = 51$) and the control group ($N = 49$)

Engagement with the Memory Task It is possible that the difference in results observed between Study 1 and Study 2 is due to the alternative manipulation of cognitive load employed. In Study 1, the control group did not

engage in any task, however, adopting De Neys and Shaeken's procedure [3], participants in the control group of Study 2 engaged in a memory task. It is possible that simply engaging in a memory task led to differences in responses, and that level of difficulty (the manipulation that was employed) was irrelevant. Indeed, the responding to the critical slide in the control group in Study 2 is more similar to the responding in the experimental group in 1 than to the control group in Study 1.

Rates of successful reproduction of the dot patterns in Study 2 were much lower than reported by De Neys and Shaeken [3]. It appears that participants in Study 2 did not engage with the memory task in the same way as participants in De Neys and Shaeken's studies [3]. We hypothesised that the effectiveness of the cognitive load manipulation may be moderated by the degree to which people engage with the manipulation.

The memory task involved correctly placing dots in a 3 x 3 grid. For the scoring of this task, each of the nine places in the grid could be marked/not marked correctly or incorrectly, making 9 the total possible number of correct responses. If a person misplaced one dot in the pattern this would count for 2 incorrect places in the grid: the mark in the incorrect place, and the absence of a mark in the place it should have been. A participant who received a score of 7, could reasonably be taken to have engaged with the task, and simply made a slip. As such, this was taken as the cut-off point for identifying engagement. This resulted in 56 participants being identified as engaging with the memory task, and 44 being identified as not engaging with the task.

Responses to critical slide for participants who engaged with the memory task and participants who did not engage with the memory task were analysed separately (see Fig. 4). A chi-squared test for independence revealed an association between engagement in the memory task and response to the critical slide, $\chi^2(2, N = 100) = 7.68$, $p = .021$, $V = .28$. The observed power was .70.

5.3   Study 2: Discussion

The aim of Study 2 was to replicate Study 1. As a replication it failed. However, interesting variability was observed when engagement with the memory task was accounted for. Two problems with the cognitive load manipulation were identified. Firstly, the control did not serve as an appropriate control. Secondly, (and unsurprisingly) the manipulation is only effective if participants engage with the memory task. A follow-up study addresses each of these.

6   Study 3 - Revised Online Replication

In Study 2 the role of engagement with the memory task emerged as an important moderator of the effectiveness of the cognitive load manipulation. Study 3 was conducted in order to test if cognitive load affects participants' ability to identify reasons for their judgements, when accounting for engagement with the memory task.

Fig. 4. Study 2: Responses to critical slide depending on engagement ($N = 56$) or non-engagement ($N = 44$) with the memory task

6.1   Study 3: Methods

Participants and Design Study 3 was a between subjects design. The dependent variable was response to the critical slide. The independent variable was cognitive load with two levels: present and absent.

Following the elimination of 34 participants who scored less than 7 on the memory task we were left with a final sample of 129 participants (74 female, 55 male; $M_{age} = 40.26$, min $= 20$, max $= 72$, $SD = 13.04$). Participants in this sample were recruited through MTurk (under the same conditions as Study 2).

Procedures and Materials Study 3 was the same as Study 2 with two changes. The control group did not take part in a memory task, and to avoid task fatigue, the dot patterns presented alternated between the easy 3-dot patterns and the complex 4-dot patterns.

A score of 7 or higher on the memory task that accompanied the critical slide was selected as the measure of engagement with the memory task. Only participants who engaged with the task were eligible for analysis. Other than the two changes described above, Study 3 was the same as Study 2.

6.2   Study 3: Results

Threre was no difference in initial judgement depending on cognitive load, $t(123.06) = 0.165$, $p = 0.870$. The responses to the critical slide for the experimental group ($N = 68$) and the control group ($N = 61$) are displayed in Fig. 5. A chi-squared test for independence revealed a significant association between experimental

condition and response to the critical slide, $\chi^2(2, N = 129) = 6.51$, $p = .039$, $V = .223$. The observed power was .63.

Fig. 5. Study 3: Responses to critical slide for the cognitive load group ($N = 68$) and the control group ($N = 61$)

6.3   Study 3: Discussion

Study 3 demonstrated the predicted relationship between engagement with a cognitive load task and providing reasons for a judgement. As expected, engagement with a cognitive load task reduced the rates of providing reasons. However, Study 3 did not include an objective manipulation check. As such a follow up study was conducted that included an objective manipulation check.

7   Study 4 - Online Replication with Manipulation Check

Participants and Design Study 4 was a between subjects design. The dependent variable was response to the critical slide. The independent variable was cognitive load with two levels: present and absent.

Following the removal of 29 participants for non-engagement with the memory task, we were left with total sample of 127 participants (84 female, 43 male; M_{age} = 41.19, min = 21, max = 74, SD = 13.91). Participants in this sample were recruited through MTurk (under the same conditions as Studies 2 and 3).

Procedures and Materials Study 4 was the same as Study 3 with one change, the inclusion of a manipulation check. A prose paragraph was included after participants made their revised judgements. Participants were then asked three comprehension questions relating to the prose paragraph. It was expected that participants in the control group would perform better at this task than participants under cognitive load [10].

Fig. 6. Study 4: Responses to critical slide for the cognitive load group ($N = 64$) and the control group ($N = 61$)

7.1   Study 4: Results

Threre was no difference in initial judgement depending on cognitive load, $t(139.46) = 0.795$, $p = 0.428$. The responses to the critical slide for the experimental group ($N = 64$) and the control group ($N = 61$) are displayed in Fig. 6. Though the pattern of responses is in the predicted direction, a chi-squared test for independence revealed no significant association between experimental condition and response to the critical slide, $\chi^2(2, N = 127) = 2.047$, $p = .359$, $V = .13$. The observed power was .23. The predicted relationship between cognitive load and dumbfounded responding was not observed in Study 4.

The responses to the manipulation check questions were investigated. There was no difference in the number of correct answers to these questions between the cognitive load group and the control group $F(1, 124) = .33$, $p = .569$, partial $\eta^2 = .003$. There was also no difference in time taken to read the vignette between the groups $F(1, 125) = 2.57$, $p = .112$, partial $\eta^2 = .020$.

7.2   Study 4: Discussion

The relationship between cognitive load and providing reasons observed in previous studies was not observed in Study 4. However, the manipulation check did not reveal any differences between the control group and the experimental group, suggesting that the cognitive load manipulation was ineffective in Study 4.

## 8   Combined Results, Discussion, and Conclusion

Fig. 7. Rates of declaring reasons and cognitive load across each study

The combined results for Studies 1-4 are displayed in Fig. 7. Two measures of engagement with the memory task were taken in the combined analysis that follows, (a) a score of 7 or higher (as in previous analyses) and (b) a score of 8 or higher. When engagement with the memory task was accounted for (using the primary measure of engagement ($N = 422$, 221 experimental, 201 control), a significant association between cognitive load and responses to the critical slide was found, $\chi^2(2, N = 422) = 12.675$, $p = .002$, $V = .17$. The observed power was .90. When the stricter measure of engagement was employed ($N = 413$, 212 experimental, 201 control), a significant association between cognitive load and response to the critical slide was found, $\chi^2(2, N = 413) = 14.847$, $p < .001$, $V = .19$. The observed power was .94. Furthermore, a mini meta-analysis was conducted and found that cognitive load significantly influenced responding across all studies in $\chi^2(8) = 23.81$, $p = .002$ (Fisher's method); or when weighting for sample size, $z = 2.95$, $p = .002$ (Stouffer's Z-score method).

The studies presented here provide some weak evidence for a conflict in dual-process explanation of moral dumbfounding (cognitive load inhibits reason

giving). The dependent variable in these studies is nominal/categorical. This is not well suited for identifying small or subtle effects. Responses in the dumbfounding paradigm are not easily influenced by experimental manipulation, leading to inconsistent results, and the findings of the studies presented here are inconclusive. Future research should employ more robust manipulations and controls for the engagement of participants, to resolve the inconsistencies in the results presented here.

References

1. Crockett, M.J.: Models of morality. Trends in Cognitive Sciences **17**(8), 363–366 (2013). https://doi.org/10.1016/j.tics.2013.06.005
2. Cushman, F.A.: Action, outcome, and value a dual-system framework for morality. Personality and Social Psychology Review **17**(3), 273–292 (2013). https://doi.org/10.1177/1088868313495594
3. De Neys, W., Schaeken, W.: When people are more logical under cognitive load: Dual task impact on scalar implicature. Experimental Psychology **54**(2), 128–133 (2007). https://doi.org/10.1027/1618-3169.54.2.128
4. Drolet, A., Frances Luce, M.: The rationalizing effects of cognitive load on emotion-based trade-off avoidance. Journal of Consumer Research **31**(1), 63–77 (2004). https://doi.org/10.1086/383424, 00152
5. Evans, J.S.B.T.: On the resolution of conflict in dual process theories of reasoning. Thinking & Reasoning **13**(4), 321–339 (2007). https://doi.org/10.1080/13546780601008825, 00173
6. Evans, J.S.B.T., Stanovich, K.E.: Dual-process theories of higher cognition: Advancing the debate. Perspectives on Psychological Science **8**(3), 223–241 (2013). https://doi.org/10.1177/1745691612460685, 00801
7. Greene, J.D.: The secret joke of kant's soul. In: Sinnott-Armstrong, W. (ed.) Moral Psychology Volume 3: The neurosciences of morality: emotion, brain disorders, and development, pp. 35–79. the MIT press (2008)
8. Haidt, J.: The emotional dog and its rational tail: A social intuitionist approach to moral judgment. Psychological Review **108**(4), 814–834 (2001). https://doi.org/10.1037/0033-295X.108.4.814
9. Haidt, J., Björklund, F., Murphy, S.: Moral dumbfounding: When intuition finds no reason. Unpublished manuscript, University of Virginia (2000)
10. Just, M.A., Carpenter, P.A.: A capacity theory of comprehension: Individual differences in working memory. Psychological Review **99**(1), 122–149 (1992). https://doi.org/10.1037/0033-295X.99.1.122
11. Kohlberg, L.: Stages in the development of moral thought and action. New York: Holt, Rinehart & Winston (1969)
12. McHugh, C., McGann, M., Igou, E.R., Kinsella, E.L.: Searching for moral dumbfounding: Identifying measurable indicators of moral dumbfounding. Collabra: Psychology **3**(1) (2017). https://doi.org/10.1525/collabra.79
13. Royzman, E.B., Kim, K., Leeman, R.F.: The curious tale of julie and mark: Unraveling the moral dumbfounding effect. Judgment and Decision Making **10**(4), 296–313 (2015)