

The case of the drunken sailor: On the generalisable wrongness of harmful transgressions

Katinka J. P. Quintelier¹, Daniel M. T. Fessler², and Delphine De Smet³

¹Department of Philosophy and Moral Sciences and Research Unit ‘The Moral Brain’, Ghent University, Ghent, Belgium

²Department of Anthropology and Center for Behavior, Evolution & Culture, University of California Los Angeles, Los Angeles, CA, USA

³Department of Legal Theory and Legal History and Research Unit ‘The Moral Brain’, Ghent University, Ghent, Belgium

There is a widespread conviction that people distinguish two kinds of acts: on the one hand, acts that are generalisably wrong because they go against universal principles of harm, justice, or rights; on the other hand, acts that are variably right or wrong depending on the social context. In this paper we criticise existing methods that measure generalisability. We report new findings indicating that a modification of generalisability measures is in order. We discuss our findings in light of recent criticisms of moral/conventional research.

Keywords: Moral/conventional distinction; Morality; Harm; Transgression; Social cognition.

Following the seminal work of Turiel (1983), scholars have held that people distinguish between what researchers label “moral transgressions” and “conventional transgressions”. While various interpretations exist, at least

Correspondence should be addressed to Katinka J. P. Quintelier, Department of Philosophy and Moral Sciences, Ghent University, Blandijnberg 2, lok. 01.28, B-9000 Ghent, Belgium. Email: katinka.quintelier@gmail.com

K. Q. received funding from the FWO (Fonds Wetenschappelijk Onderzoek) to conduct this research. The authors thank Stephen Stich for help with the pilot study, as well as two anonymous reviewers and Paulo Sousa for their comments.

two features are associated with the moral/conventional (M/C) distinction. First, individuals purportedly judge transgressions involving harm, or violations of justice or rights (hereafter: HJR transgressions), to be generalisably wrong, meaning that they are wrong across social contexts. Conversely, conventional transgressions are purportedly seen as locally wrong, meaning that people recognise that their wrongness varies across space and time, and view such variation as acceptable. Second, the distinction is characterised by the justification of the wrongness: HJR transgressions are wrong by virtue of universal principles involving harm, justice, or rights, while conventional principles are justified by convention (Turiel, 1983).

Critics have questioned the existence, exact interpretation (see Discussion section), and appropriate experimental tests of the M/C distinction (Kelly, Stich, Haley, Eng, & Fessler, 2007; Sousa, 2009; Sousa, Holbrook, & Piazza, 2009; Stich, Fessler, & Kelly, 2009). Here we raise methodological concerns regarding the measurement of generalisability.

PROBLEMS PLAGUING MEASURES OF GENERALISABILITY

According to Turiel (1983, pp. 35–36), HJR transgressions are perceived as wrong *independent of “convention”*, where Turiel (1983) understands “convention” as an implicit behavioural uniformity (e.g., practice, habit) or an explicit regulation (e.g., a rule) with a social function that can be altered by consensus or by general usage.

Previous M/C studies employ suboptimal generalisability questions. In posing the generalisability question most authors vary only one aspect of convention, leaving it to the imagination of the participant what other aspects of convention would vary or stay the same. For example, Smetana (1981) asks if it would be OK if there were no rule about it, thus only varying “explicit regulations”. Huebner, Lee, and Hauser (2010) ask if the act would still be wrong if X lived somewhere where everyone else did this; this only changes the existence of a “behavioural uniformity”. Similarly, Royzman, Leeman, and Baron (2009) vary only social norms or consensus. Nonetheless, even if a transgression is concordant with one aspect of convention (e.g., when there is no rule against it), it can still be discordant with convention by virtue of any of its other aspects (social norms, consensus, behavioural uniformity etc.).

Furthermore, in classical M/C studies the HJR transgressions are entirely different than the conventional transgressions. This—together with the suboptimal generalisability probes—creates a bias, such that generalisability questions may well invite participants to consider the conventional but not the HJR transgressions in a situation where they are in accord with

convention. For example, Huebner's generalisability question induces participants to think about a place where all children hit other children for fun (HJR transgression). However, Huebner et al. (2010) leave unspecified if, in this other place, children consent to being hit, or are also allowed by their teachers to hit other children. Participants may plausibly assume that, even if hitting is the general practice in a school in another place and time, teachers will still forbid it, or it will still be against general consensus. Thus they will generalisably condemn the HJR transgression, *not* independent of convention but because they view the situation as generally against *explicit rules* or against *implicit consensus*, i.e., against prevailing "convention". In contrast, when asked to imagine a classical "conventional" violation (e.g., addressing a teacher by the first name) happening at another place and time, participants might readily imagine that this act is not against convention elsewhere (because, for example, they know about places where this is not part of convention, leading them to be more likely to judge the act to be permissible at another time and place).

Thus existing M/C studies do not effectively test whether HJR transgressions are wrong *independent of convention*. To be effective the generalisability probe must explicitly vary every aspect of the prevailing "convention", both for the conventional *and* for the HJR transgression. Prior lack of explicitness may therefore mask a lack of symmetry across the testing conditions, which may explain the difference in response patterns in previous M/C studies.

Kelly et al. did explicitly vary several features of conventionality in a HJR transgression: They asked participants if it is wrong to whip a sailor on a modern American cargo ship, and if it was wrong 300 years ago—when it was common practice, not against the law, and deemed appropriate by everyone. Participants did indeed deem whipping more permissible in the second context. Nonetheless the scenarios might have introduced confounds having to do with harm, justice, or rights (see also Fraser, 2012): Participants might have reasoned that 300 years ago the consequences of being drunk on duty were more harmful (e.g., attacks by pirates) and therefore warranted harsher punishments, thus explaining why harsher punishments were deemed more permissible for people 300 years ago.

THE PRESENT STUDY

In this study we test whether explicitness of aspects of convention matters. We adapt Kelly et al.'s (2007) "whipping generalisability" scenario by explicitly varying one versus two aspects of convention, while avoiding the introduction of HJR confounds. We also add a justification question to check for confounds.

Hypothesis

We expect that, when few aspects of convention are explicit (i.e., when the design resembles classical M/C studies), participants will be more likely to evaluate harmful transgressions as generalisably wrong, and less likely to evaluate harmful transgressions as variably right or wrong, than when more aspects are explicit.

Design and procedure

We randomly assigned participants to one of two conditions (CONDITION: high versus low [explicitness], between-participants, 424 versus 389 participants respectively). We adapted the “whipping generalisability” scenario (see Kelly et al., 2007) such that, in the “high” condition, two aspects of convention were explicit, while in the “low” condition only one aspect was explicit; to avoid HJR confounds all other features remained the same across conditions. Each participant read two scenarios. In one scenario whipping is concordant with convention, while in the other it is against convention (SCENARIO: concordant versus discordant, within-participants). To avoid assumptions about legality we specified that whipping is legal throughout (see below for details). Pilot testing revealed that participants understood the scenarios. Order of scenarios was randomised.

After each scenario we asked participants about the permissibility of the act and their justifications for their answer. Pilot testing revealed that, when simply asked whether the act was OK, some participants made a distinction between two kinds of wrongness. They explained that it was OK legally, or that it was not against the rules, but that it was nonetheless “morally” or “ethically” wrong. We were not interested in a descriptive statement (e.g., about the legality of the act). Instead we were only interested in participants’ own evaluation of the act, and participants seemed to employ the words “morally” and “ethically” to delineate the latter. We therefore decided to ask about *moral* permissibility. We provided three forced-choice answer options: 1 = “morally permissible”, 2 = “morally permissible but wrong for reasons that have nothing to do with morality (e.g., it might be unlawful)”, and 3 = “morally wrong”.

The “high” scenarios are presented below, with relevant differences between them underlined for clarity; italicised sections were omitted in the “low” condition. Participants saw the scenarios in plain text.

Discordant (with convention)

Mr Johnson is an officer on a cargo ship in 2010, carrying goods along the Atlantic coastline. All the crew members are American but the ship is mostly

in international waters. When a ship is in international waters, it has to follow the law of the state whose flag it sails under and each ship can sail under only one flag. This ship does not sail under the U.S. flag. The law of this ship's flag state allows both whipping and food deprivation as a punishment.

On this ship, food deprivation is always used to discipline sailors who disobey orders or who are drunk on duty; *as a consequence everyone on this ship has come to think that food deprivation is an appropriate punishment.* Whipping however is never used to discipline sailors *and no one on this ship thinks whipping is an appropriate punishment.*

One night, while the ship is in international waters, Mr Johnson finds a sailor drunk at a time when the sailor should have been on watch. After the sailor sobers up, Mr Johnson punishes the sailor by giving him 5 lashes with a whip. This does not go against the law of the flag state.

Participants answered the following permissibility question:

Is it morally permissible for Mr Johnson to whip the sailor? (multiple-choice format)

1. Yes, it is morally permissible
2. Yes, it is morally permissible but it is wrong for reasons that have nothing to do with morality (e.g., it might be unlawful)
3. No, it is morally wrong (whether it is right or wrong in other ways or not)

Depending on their answer, participants were then led to the appropriate justification question(s):

(1 or 2) → Why is it morally permissible for Mr Williams to whip the sailor? Please briefly list all the reasons you thought of while you were answering question 1. (Write a minimum of one word.)

(2) → In what other ways is it wrong for Mr Williams to whip the sailor? Please briefly list [...]

(3) → Why is it morally wrong for Mr Williams to whip the sailor? Please briefly list [...]

Concordant (with convention)

Mr Williams is an officer on another cargo ship in 2010, [...]

On this ship, whipping is always used to discipline sailors who disobey orders or who are drunk on duty; *as a consequence everyone on this ship has come to think that whipping is an appropriate punishment.* Food deprivation however is never used to discipline sailors *and no one on this ship thinks food deprivation is an appropriate punishment.*

One night, while the ship is in international waters, Mr Williams finds a sailor drunk at a time when the sailor should have been on watch. After the sailor

sobers up, Mr Johnson punishes the sailor by giving him 5 lashes with a whip. This does not go against the law of the flag state.

Participants again answered the permissibility and justification questions.

Participants

This study was deemed exempt from ethical review by the UCLA Institutional Review Board. We recruited participants using Amazon.com's Mechanical Turk web-based employment system. Discarding incomplete responses produced a sample of 1024 participants (536 men; $M_{age} = 30.83 \pm 10.56$ *SD*; range = 18–78). Participants indicated the following countries of residence: USA (492), India (388), other (144). Discarding inconsistent responses (see below) produced a final sample of 813 participants (408 men; $M_{age} = 31.13 \pm 10.72$ *SD*; range = 18–78; USA: 415, India: 282, elsewhere: 115).

Analyses

Permissibility question. We expected *a priori* that answer option 2 (“morally permissible but wrong for other reasons”) could be pooled with answer option 1 (“morally permissible”). However, a close reading of participants' justifications¹ revealed that this did not necessarily accord with participants' interpretation of the question. For example, one participant gave the following justifications after choosing option 2:

Why is it morally permissible? “Because it is permitted by law”
 Why is it wrong? “Because it is not used on this ship. Sailors should know what punishment they should expect.”

In this example it is not clear which of the two justifications provided is more evaluative, making it impossible to determine whether the response is closer to option 1 or option 3. On the basis of the corresponding justifications, answers to option 2 could not clearly be pooled with either option 1 or 3 in 253 of 278 cases (91% of option 2 justifications; first author's categorisation). We therefore pooled those 25 clear answers with option 1 or 3, as appropriate, and removed the 187 participants who, in response to at least one of the vignettes, had chosen answer option 2 and

¹ Two reviewers asked whether pooling answer option 2 with answer option 1 had an effect on the results. We therefore explored whether employing different inclusion criteria (based on the justifications) altered the results; however, this procedure did not affect the results (see Results section).

provided a justification that could not be pooled with 1 or 2. We thus retained 813 participants.

Response patterns. To evaluate the extent to which participants deemed the acts to be generalisably wrong, we created four categories based on their answers to the permissibility questions, thus creating an additional variable CATEGORY (see Table 1). If participants answered that whipping was permissible in both scenarios—cf. YES–YES in Kelly et al. (2007) and Sousa (2009)—their answers are categorised as “generalisably permissible”. If they answered it is wrong in both scenarios (NO–NO), their answer falls under the category “generalisably wrong”. Their answers were classified as “variably right or wrong” when they deemed whipping to be wrong when against convention but permissible when in accord with convention (NO–YES). The remainder (YES–NO) were classified as “unexpected”, as we do not have a ready explanation for these answer patterns beyond the possibility of participant confusion. Percentages and frequencies are listed later in Table 3.

Justification. In order to control for HJR confounds, two independent coders categorised participants’ justifications. Following M/C theory, if a participant referred to universal principles of harm, justice or rights, this was categorised as a HJR justification. If a participant referred to consensus, consent, explicit rule, law, or practice, this was categorised as a context justification. Inter-coder agreement for all justifications was moderate (Cohen’s kappa = .561; $n = 4096$); however, the results were the same no matter which coder we used. We therefore report the results for the mean of both coders later in Table 3.

TABLE 1
% (#) of participants morally permitting and not permitting the transgression
(per scenario and condition)

	<i>High explicitness condition</i>		<i>Low explicitness condition</i>	
	<i>Discordant scenario</i>	<i>Concordant scenario</i>	<i>Discordant scenario</i>	<i>Concordant scenario</i>
Morally permissible	35.0 (158)	52.2 (247)	30.3 (125)	44.7 (199)
Not morally permissible	65.0 (293)	47.8 (226)	69.7 (287)	55.3 (246)
TOTAL	100 (451)	100 (473)	100 (412)	100 (445)

The number of participants differs between scenarios within condition. At first pass this may seem odd given that this was a within-participant factor. However, this occurs because we either removed answer option 2 when it occurred or pooled it with answer option 1 or 3 where appropriate (see Permissibility section). Some participants gave answer option 2 in response to one scenario, but not in response to the other scenario.

Results

The descriptive and inferential statistics for the full design (mixed-model repeated-measures ANOVA: CONDITION [between-participants] \times SCENARIO [within-participants]) are presented in Tables 1 and 2.

The results relevant to our hypothesis (“Is there a relationship between CONDITION and CATEGORY?” see earlier Hypothesis section) are presented in Table 3.

Chi square tests revealed that order had no effect on categorisation, CATEGORY: $\chi^2(3, 813) = 2.508$; $p = .474$. We therefore pooled the two orders, finding a general effect of CONDITION on the distribution of categories: $\chi^2(3, 813) = 11.106$; $p < .05$. However, the effect size of CONDITION on CATEGORY was small (*Cramer’s V* = .117). The adjusted residuals (Haberman, 1978, p. 197) revealed that participants assigned to the “high” condition were significantly less likely to judge the act to be generalisably wrong ($p < .05$), and were also significantly more likely to judge the act to be variably right or wrong, than participants assigned to the other condition ($p < .05$).

Additional analysis revealed that removing the 187 participants had not substantially changed the results. When pooling all answer options 2 with 1, as initially planned, we again found a general effect of CONDITION on distribution of categories: $\chi^2(3, 1021) = 8.096$; $p < .05$. The effect size of CONDITION on CATEGORY was again very small (*Cramer’s V* = .064). Nevertheless, the adjusted residuals again revealed that participants assigned to the high-explicitness condition were less likely to judge the act to be generalisably wrong ($p < .05$), and were also again more likely to judge the act to be variably right or wrong, than participants assigned to the other condition ($p < .05$). The same pattern of results appears if we simply remove all participants who had chosen answer option 2, thus retaining 800 participants.

If HJR confounds caused the variable wrongness, participants would refer to universal principles of harm, justice, or rights in justifying their evaluations of permissibility. However, of the 103 participants displaying conventional response patterns, only 6 referred to HJR principles when justifying why they thought the whipping in accord with convention was

TABLE 2
Inferential statistics for mixed-model repeated-measures ANOVA

	<i>F</i>	<i>p</i>	<i>Partial eta squared</i>
CONDITION (between)	4.245 (1, 811)	< .05	.005
SCENARIO (within)	77.507 (1, 811)	< .001	.087
CONDITION \times SCENARIO	4.260 (1, 811)	< .05	.005

TABLE 3
Distribution of participants and their justifications according to CATEGORY

CATEGORY	Participants per CATEGORY × CONDITION % (#)		Mean % (#) of participants giving this kind of justification per category				TOTAL % (#) participants per category
	High	Low	Discordant with convention		Concordant with convention		
			HJR	Context	HJR	Context	
Generalisably permissible yes-yes	32.3 (137)	29.3 (114)	16.7 (42)	72.1 (181)	15.1 (38)	69.7 (175)	30.9 (251)
Generalisably wrong no-no	50.2 (213)	60.2 (234)	81.5 (365)	15.5 (70)	81.3 (364)	11.6 (52)	55 (447)
Variably right or wrong no-yes	15.6 (66)	9.5 (37)	22.8 (24)	72.8 (75)	5.8 (6)	85.9 (89)	12.7 (103)
Unexpected yes-no	1.9 (8)	1.0 (4)	16.7 (2)	79.1 (10)	50.0 (6)	33.3 (4)	1.5 (12)
	100 (424)	100 (389)					100 (813)

morally permissible. In comparison, 364 of the 447 participants who gave a moral response pattern referred to HJR principles to justify why they deemed whipping in accord with convention to be wrong.

Discussion

Results provided partial support for our contention that measurement bias might have contributed to the frequently reported finding that HJR transgressions are widely viewed as generalisably wrong: On the one hand, our participants were significantly more likely to assert that hypothetical HJR transgressions are generalisably wrong when they were provided with less information regarding the relevant conventions than when more aspects of convention explicitly permitted the behaviour. This suggests that, in prior research, participants who were thought to be asserting generalisability might, in fact, have merely been inferring the uniformity of conventions across space and time. On the other hand, effect sizes were admittedly quite small. Unfortunately, to the best of our knowledge, effect sizes have not previously been reported in the M/C literature. This makes it difficult to assess the extent to which measurement bias can explain the effects in M/C studies. Moreover, it is important to note that we varied only two versus one aspect(s) of convention. It would be interesting to investigate whether, as our position entails, the effect of explicitness on response patterns is greater when more than two aspects of convention (versus only one) are made explicit.

Sousa (2009) criticises permissibility questions such as ours because they are polysemous. In asking “Is it OK for X to H?” researchers have an evaluative question in mind. However, participants can also interpret this in a descriptive way, i.e., “According to X, is it OK to H?”

Our permissibility questions are not immune to this criticism. However, participants’ justification answers can be used to check for these confounds. Indeed, we found that some participants’ justifications revealed a descriptive reading of the permissibility question, as the following example illustrates:

Why is it morally wrong? “It was made clear that it was lawfully acceptable for the sailor to be whipped; *but in my mind what is morally right and what is lawfully right are not the same.*” [our italics.]

The non-italicised part points to a descriptive interpretation. Nonetheless, this did not drive the participant’s permissibility answer: consonant with the italicised evaluative reading, the participant answered that the (concordant) act was morally wrong. Thus this was not a reason to re-categorise the participant’s answer.

In general, clear descriptive justifications were rare, and when they occurred they did not necessitate discarding the participant or re-categorising the

permissibility answer. While descriptive readings did not appear to pose a problem for our probes, we nonetheless agree with Sousa (2009) that investigators must be aware of this possibility. We also agree that justification questions are one possible way to check for confounds. However, it would be better to develop non-ambiguous permissibility questions in addition to multiple-choice questions that allow for a more efficient search for descriptive readings.

How does our study speak to the M/C distinction? It all depends on how one interprets the M/C distinction. Sousa (2009) offers a “deflationary hypothesis”, arguing that harmful transgressions (such as ours) are not appropriate for testing the M/C distinction. Instead, investigators should use transgressions that are harmful *and* violate principles of justice or rights. Sousa adduces support for this view from several sources. First, a harmful act is not necessarily a transgression. The M/C distinction holds that only harmful *transgressions* will be deemed generalisably wrong, and this condition, *dixit* Sousa (2009), is satisfied if the act is also an injustice or a rights violation. Second, Sousa claims that prevailing practice supports his view. He argues that, to test if the act is really a transgression, participants should judge it to be wrong in the discordant scenario; participants who do not do so should be excluded from the analysis (Sousa notes that some researchers do exactly this). Third, Sousa reanalyses Kelly et al.’s findings and, after removing putatively “irrelevant” participants, finds that the frequency of “generalisably wrong” answer patterns is very high, supporting his claim that Kelly et al. (2007) did not find convincing evidence against the M/C distinction.

We disagree with Sousa’s (2009) argument for three reasons. First, Turiel (1983, pp. 34–46) provides no indication that harm is not enough to elicit the moral response pattern. On the contrary, he argues that children learn to associate harm with undesirable consequences without knowledge of the institutional context: they merely need to see harm to empathise and, as a consequence, learn that harm is intrinsically bad. (Turiel, 1983, pp. 39–40, argues that similar things occur in the case of violations of justice.) Thus the original thesis is that harmful acts (just as violations of justice) in themselves elicit the “generalisably wrong” response pattern.

Second, this does not preclude adopting Sousa’s novel reading of the theory, but said interpretation must be clarified if it is to not produce a self-fulfilling prediction. While some rights are perceived as universal, others are not (see also Stich et al., 2009). If Sousa (2009) were to limit the category of HJR transgressions to harmful acts that also violate “rights that are perceived to be universal”, then HJR transgressions will by definition be held to be generalisably wrong. Conversely, if Sousa were to include in said criteria rights that are perceived to depend on the social context then, by Turiel’s definitions, said rights will be matters of convention and, as a consequence, they will not be perceived as generalisably wrong independent

of convention. Proponents of Sousa's (2009) position might counter that the *idea* of justice or rights is perceived as universal, yet specific principles can be perceived to vary. However, this does not help, as the *idea* of convention may also be perceived as universal while its specific rules are perceived to vary. Hence it is not clear how "justice and rights" are to be conceptualised such that they differ from convention in a non-circular way (see also Stich et al., 2009).

Third, in light of the above, we see no need to remove so-called "irrelevant" participants. Nevertheless, the most effective way of precluding Sousa's critique is to simply employ Sousa's prescription. After removing participants whom Sousa deems "irrelevant," we find 76% (213/66+213) "generalisably wrong" answer patterns in the high-explicitness condition, and 86% (234/37+234) in the low-explicitness condition. These fractions are indeed higher than occur in the total sample (51% and 60%, high versus low, respectively). However, the percentage of "variably right or wrong" answer patterns is also higher after such filtering: 23% (high-explicitness) and 14% (low-explicitness) versus the previous 15% and 10%, respectively. Revisiting the data presented in Sousa (2009) and Kelly et al. (2007) reveals that Sousa's filtering has similar effects in these datasets as well. In short, Sousa's prescription makes no difference.

What, then, can we infer from a minority response? Again, it all depends on how one interprets the M/C distinction. Arguably, a central pillar of the M/C distinction is that HJR transgressions uniformly evoke the "generalisably wrong" answer pattern. We find that this is not so: a substantial minority of participants is divided among the categories "generalisably permissible" and "variably right or wrong". We found a lower percentage of the latter than did Kelly et al. (2007) and Sousa (2009), who used a similar, albeit confounded, scenario. HJR confounds or descriptive readings of the permissibility question may have contributed to this. However, as Kelly et al. (2007) did not collect justifications, and as we employed a more culturally diverse pool of participants, it is difficult to determine what is responsible for the differences across studies.

A weaker interpretation of the M/C distinction is that HJR transgressions are more likely to elicit "generalisably wrong" answer patterns, and less likely to elicit "variably right or wrong" answer patterns, than are conventional transgressions. Our study does not directly speak to this. Sceptics may therefore argue that our results are not informative because we did not compare whipping with non-harmful transgressions—perhaps our percentages would be dwarfed by the fraction of participants that would deem "conventional" transgressions as variably wrong. However, we suggest that the outcome of such a comparison would importantly hinge on the "conventional" transgressions employed. Huebner et al. (2010) document the existence of a range of acts that do not reliably elicit either

“variably right or wrong” or “generalisably wrong” answer patterns. Nichols (2002) found that disgusting violations that would otherwise be classified as conventional were deemed less permissible and more seriously wrong than non-disgusting conventional transgressions. Therefore whether conventional transgressions would elicit a smaller proportion of generalised responses than did the transgression we employed would likely depend on what conventional issues were chosen.

To conclude, the literature on the M/C distinction is fraught with confusion. The present study suggests that the oft-reported finding that HJR transgressions are widely viewed as generalisably wrong may well be at least in part a methodological artefact. Attempts to defend the Turiel position suffer from inadequately specified concepts, and are at risk of inherently precluding disconfirmation. The onus is therefore on proponents of the Turiel position to reconfigure their methods and revisit their core concepts.

Manuscript received 29 July 2011

Revised manuscript received 7 February 2012

First published online 11 May 2012

REFERENCES

- Fraser, B. (2012). The nature of moral judgments and the extent of the moral domain. *Philosophical Explorations: An International Journal for the Philosophy of Mind and Action*, 15(1), 1–16.
- Haberman, S. J. (1973). The analysis of residuals in cross-classified tables. *Biometrics* 29(1): 205–220. Stable URL: <http://www.jstor.org/stable/2529686>.
- Haberman, S.J. 1978. *Analysis of qualitative data*. London: Academic Press.
- Huebner, B., Lee, J. J., & Hauser, M. D. (2010). The moral–‘conventional’ distinction in mature moral competence. *Journal of Cognition and Culture*, 10, 1–26. doi: 10.1163/156853710x497149.
- Kelly, D., Stich, S., Haley, K. J., Eng, S. J., & Fessler, D. M. T. (2007). Harm, affect, and the M/C distinction. *Mind & Language*, 22(2), 117–131. doi: 10.1111/j.1468-0017.2007.00302.x.
- Nichols, S. (2002). Norms with feeling: Towards a psychological account of moral judgment. *Cognition*, 84(2), 221–236. doi: 10.1016/S0010-0277(02)00048–3.
- Royzman, E. B., Leeman, R. F., & Baron, J. (2009). Unsentimental ethics: Towards a content-specific account of the moral–‘conventional’ distinction. *Cognition*, 112(1), 159–174. doi: 10.1016/j.cognition.2009.04.004.
- Smetana, J. G. (1981). Preschool childrens conceptions of moral and social rules. *Child Development*, 52(4), 1333–1336.
- Sousa, P. (2009). On testing the ‘moral law’. *Mind & Language*, 24(2), 209–234. doi: 10.1111/j.1468-0017.2008.01360.x.
- Sousa, P., Holbrook, C., & Piazza, J. (2009). The morality of harm. *Cognition*, 113(1), 80–92. doi: 10.1016/j.cognition.2009.06.015.
- Stich, S., Fessler, D. M. T., & Kelly, D. (2009). On the morality of harm: A response to Sousa, Holbrook and Piazza. *Cognition*, 113(1), 93–97. doi: 10.1016/j.cognition.2009.06.013.
- Turiel, E. (1983). *The development of social knowledge. Morality and social ‘convention’*. Cambridge, UK: Cambridge University Press.

Copyright of Thinking & Reasoning is the property of Psychology Press (UK) and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.