**Not as distinct as you think: Reasons to doubt that morality comprises a unified and
objective conceptual category**

Jordan Theriault^a and Liane Young^b

<http://dx.doi.org/10.1017/S0140525X18000195>

^aDepartment of Psychology,
Northeastern University,
Boston, MA 02115

^bDepartment of Psychology,
Boston College,
Chestnut Hill, MA 02467

Commentary on:

Stanford, P. (2018). The difference between ice cream and Nazis: Moral externalization and the evolution of human cooperation. *Behavioral and Brain Sciences*, 41, E95.
doi:10.1017/S0140525X17001911

Author response:

Stanford, P. (2018). Moral externalization and normativity: The errors of our ways. *Behavioral and Brain Sciences*, 41, E119. doi:10.1017/S0140525X17002254

Abstract: That morality comprises a distinct and objective conceptual category is a critical claim for the target article. We dispute this claim. Statistical conclusions about a distinct moral domain were not justified in prior work, on account of the “stimuli-as-fixed-effects” fallacy. Furthermore, we have found that, behaviorally and neurally, morals share more in common with preferences than facts.

In the target article, Stanford argues that moral demands inhabit a distinct conceptual category, where they are experienced as externally imposed obligations, and that evolutionarily, this externalization protected prosocial individuals from exploitation, ensuring that any felt obligation to conform to a social norm was paired with a conviction that others should conform as well. Hence, externalizing moral demands (i.e., experiencing moral demands as objective) allowed individuals to reap the benefits of prosociality while also policing defectors.

One critical claim for this argument is that morality comprises “a distinctive conceptual category” (Stanford, sect. 5, para. 15, sect 6, para. 1). Work was reviewed, showing that children categorically distinguish morals from social conventions (Smetana 2006; Turiel 1983); moral properties are distinguished from response-dependent properties (e.g., “yucky”; Nichols & Folds-Bennett 2003); and, critically, morals are rated as categorically more objective than preferences and social conventions (Goodwin & Darley 2008; 2012), licensing the conclusion that “[moral beliefs are] treated almost as objectively as scientific or factual beliefs ... [and] as categorically different from social conventions (Goodwin & Darley 2008, p. 1359). However, we doubt that morals comprise a distinct category – at least on the basis of their objectivity, universality, and authority independence, as has traditionally been argued. First, prior work has not licensed statistical generalizations about a moral domain, as its authors had assumed. Second,

our recent work suggests that objectivity is not an essential feature of morality; rather, behaviorally and neurally, moral claims are more akin to preferences.

To preface our statistical criticism: in most cases, morality must be studied using specific stimuli. For example, stimuli might include asking children whether hitting (e.g., Wainryb et al. 2004) or stealing (Tisak & Turiel 1988) is acceptable. Goodwin and Darley (2008) asked about discrimination, robbery, and firing into a crowd, among others. The statistical problem is that one must move, by inference, from the specific stimuli to a sampled population (e.g., a moral domain). To make this inference, stimuli must be treated as a *random effect* (i.e., as a random sample from a population). If stimuli are averaged, then statistical conclusions (e.g., a *t*-test across subjects) apply only to those stimuli; in this case, stimuli are a *fixed effect* (and this leaves aside issues with randomly sampling moral stimuli, a problem beyond the scope of this commentary). This “stimuli-as-fixed-effects” fallacy has been identified in other fields (Clark 1973) and is easily solved using mixed effects analyses (Baayen et al. 2008), but the problem has been largely ignored within psychology (Judd et al. 2012; Westfall et al. 2014). The moral/conventional distinction has been criticized on the basis of stimulus content (e.g., stimuli typically describe “schoolyard” violations; Kelly et al. 2007, p. 121), and these criticisms may be justified, but ultimately, such content-based criticisms are unnecessary, as the original findings never licensed conclusions about a moral domain at all. They licensed conclusions about the *exact* stimuli that were used.

Prior work has argued that morality is essentially objective, but how to measure meta-ethical judgment has also been a long-standing concern (for an excellent discussion, see Goodwin & Darley 2010). Stanford rightly calls attention to the “hybrid character” of morality, where moral claims fall somewhere between “[objective] representations of how things stand in

the world itself ... and our subjective reactions to those states of the world” (sect. 6, para. 11); however, prior work has rarely allowed participants to express this hybrid nature. For example, if participants are forced to classify moral claims as true, false, or an opinion/attitude (Goodwin & Darley 2008), then distinctions between morals, facts, and preferences may appear to be more discrete than they actually are. We attempted to address this issue in a recent study (Theriault et al., 2017), where participants read moral claims (presented alongside facts and preferences) and simultaneously rated the extent that each was “about facts,” “about preferences,” and “about morality” (1–7; “not at all” to “completely”). Moral claims should be more moral-like than fact-like or preference-like; however, the question of interest was which secondary feature would dominate: Are morals largely fact-like? Or are they largely preference-like?

Although prior work has emphasized that moral claims are essentially objective, our work suggested the opposite: that moral claims were perceived as largely preference-like. Among a set of 24 moral claims that we had generated, and also among 22 claims adapted from the moral foundations questionnaire (Graham et al. 2011; Iyer et al. 2012), participants rated moral claims as significantly more preference-like than fact-like. Furthermore, we scanned subjects as they read the same claims and found that moral claims elicited widespread activity in brain regions for social cognition and theory of mind (Schurz et al. 2014; Van Overwalle 2009), overlapping with activity for preferences, but not facts. Stanford argued that humans have “[gone] in for cognitively complex forms of representation,” and that moral norms have likely been shoehorned into an evolved framework where “the most fundamental division ... [is] between how things stand in the world ... and our subjective reactions to those states” (sect. 6, para. 11). If this fundamental representational division exists, and if moral demands were (in part) externalized by coopting cognitive processes that evolved to represent the world, as

Stanford seems to argue, then we should see at least some significant overlap between processing for morals and facts. Instead, morals were behaviorally perceived, and neurally represented, as akin to preferences.

Nevertheless, we agree that moral demands are often experienced as external, even if the moral domain is not as unified as prior work has suggested. But this moral externalization may exist along a spectrum: Some moral claims may be experienced as more objective than others. Indeed, we characterized this variability in a recent study, where by-stimuli moral objectivity tracked with activity in social brain regions (Theriault et al., under review). Understanding this variability will be critical for an account of why some moral demands are experienced as obligatory and enforced (e.g., “murder is wrong”) whereas others are not (e.g., “eating meat is wrong”).

References

- Baayen, R. H., Davidson, D. J. & Bates, D. M. (2008) Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59:390–412. <http://dx.doi.org/10.1016/j.jml.2007.12.005>.
- Clark, H. (1973) The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior* 12:335–59. [http://dx.doi.org/10.1016/s0022-5371\(73\)80014-3](http://dx.doi.org/10.1016/s0022-5371(73)80014-3).
- Goodwin, G. P. & Darley, J. M. (2008) The psychology of meta-ethics: Exploring objectivism. *Cognition* 106:1339–66. <http://dx.doi.org/10.1016/j.cognition.2007.06.007>.
- Goodwin, G. P. & Darley, J. M. (2010) The perceived objectivity of ethical beliefs: Psychological findings and implications for public policy. *Review of Philosophy and Psychology* 1:161–88. <http://dx.doi.org/10.1007/s13164-009-0013-4>.
- Goodwin, G. P. & Darley, J. M. (2012) Why are some moral beliefs perceived to be more objective than others? *Journal of Experimental Social Psychology* 48:250–56. <http://dx.doi.org/10.1016/j.jesp.2011.08.006>.
- Graham, J., Nosek, B. A., Haidt, J., Iyer, R., Koleva, S. & Ditto, P. H. (2011) Mapping the moral domain. *Journal of Personality and Social Psychology* 101(2):366–85. <http://dx.doi.org/10.1037/a0021847>.
- Iyer, R., Koleva, S., Graham, J., Ditto, P. & Haidt, J. (2012) Understanding libertarian morality: The psychological dispositions of self-identified libertarians. *PLoS ONE* 7(8):e42366. <http://dx.doi.org/10.1371/journal.pone.0042366>.
- Judd, C. M., Westfall, J. & Kenny, D. A. (2012) Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored

- problem. *Journal of Personality and Social Psychology* 103:54–69.
<http://dx.doi.org/10.1037/a0028347>.
- Kelly, D., Stich, S., Haley, K. J., Eng, S. J. & Fessler, D. M. T. (2007) Harm, affect, and the moral/conventional distinction. *Mind & Language* 22:117–31.
<http://dx.doi.org/10.1093/acprof:oso/9780199733477.003.0013>.
- Nichols, S. & Folds-Bennett, T. (2003) Are children moral objectivists? Children's judgments about moral and response-dependent properties. *Cognition* 90:B23–32.
[http://dx.doi.org/10.1016/s0010-0277\(03\)00160-4](http://dx.doi.org/10.1016/s0010-0277(03)00160-4).
- Schurz, M., Radua, J., Aichhorn, M., Richlan, F. & Perner, J. (2014) Fractionating theory of mind: A meta-analysis of functional brain imaging studies. *Neuroscience and Biobehavioral Reviews* 42:9–34. <http://dx.doi.org/10.1016/j.neubiorev.2014.01.009>.
- Smetana, J. (2006) Social-cognitive domain theory: consistencies and variations in children's moral and social judgments. In: *Handbook of moral development*, ed. M. Killen & J. Smetana, pp. 119–153. Lawrence Erlbaum Associates.
- Theriault, J., Waytz, A., Heiphetz, L. & Young, L. (2017) Examining overlap in behavioral and neural representations of morals, facts, and preferences. *Journal of Experimental Psychology: General*, 146(3), 305–317. <http://dx.doi.org/10.1037/xge0000350>.
- Theriault, J., Waytz, A., Heiphetz, L. & Young, L. (under review) Theory of mind network activity is associated with metaethical judgment: An item analysis. *PsyArXiv*,
<http://dx.doi.org/10.17605/OSF.IO/GB5AM>.
- Tisak, M. S. & Turiel, E. (1988) Variation in seriousness of transgressions and children's moral and conventional concepts. *Developmental Psychology* 24:352–57.
<http://dx.doi.org/10.1037/0012-1649.24.3.352I>.

- Turiel, E. (1983) *The development of social knowledge*. Cambridge University Press.
- Van Overwalle, F. (2009) Social cognition and the brain: A meta-analysis. *Human Brain Mapping* 30:829–58. <http://dx.doi.org/10.1002/hbm.20547>.
- Wainryb, C., Shaw, L. S., Langley, M., Cottam, K. & Lewis, R. (2004) Children's thinking about diversity of belief in the early school years: Judgments of relativism, tolerance, and disagreeing persons. *Child Development* 75:687–703. <http://dx.doi.org/10.1111/j.1467-8624.2004.00701.x>.
- Westfall, J., Kenny, D. A. & Judd, C. M. (2014) Statistical power and optimal design in experiments in which samples of participant respond to samples of stimuli. *Journal of Experimental Psychology: General* 143:2020–45. <http://dx.doi.org/10.1037/xge0000014>.