



Philosophy instruction changes views on moral controversies by decreasing reliance on intuition

Kerem Oktar^{a,*}, Adam Lerner^b, Maya Malaviya^a, Tania Lombrozo^a

^a Department of Psychology, Princeton University, United States of America

^b Department of Philosophy, Princeton University, United States of America

ARTICLE INFO

Keywords:

Morality
Ethics
Belief
Controversy
Intuition
Deliberation

ABSTRACT

What changes people's judgments on moral issues, such as the ethics of abortion or eating meat? On some views, moral judgments result from deliberation, such that reasons and reasoning should be primary drivers of moral change. On other views, moral judgments reflect intuition, with reasons offered as post-hoc rationalizations. We test predictions of these accounts by investigating whether exposure to a moral philosophy course (vs. control courses) changes moral judgments, and if so, via what mechanism(s). In line with deliberative accounts of morality, we find that exposure to moral philosophy changes moral views. In line with intuitionist accounts, we find that the mechanism of change is reduced reliance on intuition, not increased reliance on deliberation; in fact, deliberation is related to increased confidence in judgments, not change. These findings suggest a new way to reconcile deliberative and intuitionist accounts: Exposure to reasons and evidence can change moral views, but primarily by discounting intuitions.

Philosophers have long argued that moral deliberation is key to fostering change in people's notions of right and wrong (e.g., Parfit, 2011). Yet psychological evidence regarding the efficacy of moral deliberation is mixed: Some studies suggest that deliberation can influence views (e.g., Horne, Powell, & Hummel, 2015), others indicate that moral judgment is driven by intuition instead (e.g., Herec et al., 2022). Though much research in moral psychology supports the intuitionist perspective (Haidt, 2001), recent meta-analyses in applied psychology show that putatively deliberative interventions—in the form of ethics courses—systematically alter moral judgments (e.g., Watts et al., 2017). These findings are difficult to reconcile, in part because the mechanisms through which ethics instruction influences moral views has not been effectively investigated. In this paper, we aim to address this gap through an experiment that jointly examines the plasticity of moral views in response to moral education—as well as the psychological basis of such plasticity—across many important real-life controversies.

Prior work on moral development and moral education has documented the effects of extended instruction on moral judgments (see Nucci, Narvaez, & Krettenauer, 2014). For instance, studies of business and science ethics courses have examined the impact of these courses on awareness of domain-specific ethical issues (e.g., conflicts of interest in

business; Ritter, 2006), judgments in canonical moral dilemmas (e.g., whether the hungry are permitted to steal from the rich; Fletcher-Brown, Buono, Frederick, Hall, & Sultan, 2012), and attitudes about domain-specific moral issues (e.g., conducting research without IRB approval; Mumford et al., 2008). Meta-analyses of these studies show small-to-moderate effects across measures for business ethics courses (overall $d = 0.29$ from 38 studies; Waples, Antes, Murphy, Connelly, & Mumford, 2009) and science ethics courses (overall $d = 0.48$ from 150 studies; Watts et al., 2017), such that enrolled students were increasingly aware of ethical issues, reflective in their responses to dilemmas, and likely to endorse appropriate domain-specific judgments (see also Schwitzgebel, Cokelet, & Singer, 2020).

If these presumably deliberative interventions change moral views, why have studies in moral psychology failed to find effects of deliberation on moral judgments, and instead supported intuitionist perspectives? One explanation is that moral deliberation may not have been given a fair shot (Bloom, 2010): Attentively examining evidence and arguments requires time, intent, and effort—but past studies have often used minimal interventions (e.g., completing a three-item inventory, Paxton, Ungar, & Greene, 2012), short timespans (e.g., minutes to hours, Horne et al., 2015), and artificial scenarios (e.g., trolley problems, Bago et al., 2022).

* Corresponding author at: Dept. of Psychology, Princeton University, Princeton, NJ 08540, United States of America.

E-mail address: oktar@princeton.edu (K. Oktar).

In this paper, we report a study that addresses this limitation of past work by investigating moral view change and its psychological basis in a setting optimal for deliberation: an introductory moral philosophy course. We find that taking this course leads to substantial changes in views across key controversies, decreases reliance on intuition, and increases reliance on deliberation in evaluating moral claims. Importantly, decreased reliance on intuition predicted *changes* in moral views, whereas increased reliance on deliberation predicted increased *confidence* in moral judgments.

1. Experiment

We presented 14 moral controversies to undergraduates enrolled in an introductory moral philosophy course at Princeton University. We measured students' views about the controversies at the beginning and end of the semester. We also measured the extent to which their views were based on intuition and deliberation. To ensure that any effects were specifically the consequence of taking moral philosophy (vs. higher-education more broadly), we collected data from a control group—a pool of students not enrolled in the moral philosophy course—over the same time period.

The study was pre-registered, and all key comparisons were conducted as pre-registered (see our repository for pre-registration, materials, data; <https://osf.io/y5tdu/>).

2. Methods

2.1. Participants

Participants were 199 adults: 137 volunteers from the moral philosophy course (79 male, 53 female, 5 other, mean age = 20) and 62 recruited from the Princeton Student Subject Pool (19 male, 41 female, 2 other, mean age = 20) in exchange for participation credit. We did not base our sample on a power analysis, as we did not have a priori expectations for effect sizes. We instead aimed to recruit as many students as possible from both the philosophy course (henceforth PHI) and the subject pool. Subject pool participants were all enrolled in at least one psychology course during the semester, and were excluded from the control group if they were also enrolled in PHI (this resulted in one exclusion). An additional 10 participants from the philosophy course and 18 participants from the control group were not included for not completing both the pre- and post-test. This experiment was approved by the Princeton Institutional Review Board (IRB).

2.2. Materials and procedure

Participants received 14 moral controversies in random order. We chose the controversies to overlap with content from PHI and to target issues actively debated within moral philosophy (see our repository for course materials). Two of these controversies (on seatbelt laws and vaccine mandates) were not addressed in the course (in discussion or readings) due to time limitations, and no change in judgments were observed. We therefore exclude these items from analyses (their inclusion does not change our results; analyses available in our repository).

Participants were asked, “To what extent do you find [controversial action] morally acceptable or morally wrong?”, with the action filled in with each controversial action presented in the first column of Table 1. Judgments were elicited on 7-point scales from “Morally wrong in most or all cases” (1) to “Morally right in most or all cases” (7), with “Neither morally acceptable nor wrong” as a neutral midpoint (4).

After each view judgment, participants answered the following meta-cognitive items in random order, measured on 7-point scales from “Not at all” (1) to “Entirely” (7):

[Confidence:] How confident are you in this moral judgment?

[Deliberation:] Is your judgment based on deliberation/analysis?

[Intuition:] Is your judgment based on intuition/emotion?

Table 1

Table of moral controversies used in the study.

Controversial action	Elaboration of relevant moral issues	Relevant references
Eating meat	Is it morally permissible to raise animals and kill them to eat their bodies?	Cohen, 1986; Norcross, 2004.
Buying fancy stuff (e.g., luxury cars, big homes) instead of donating to charity	Is it morally permissible to spend money on non-essential goods when that money could be used to save or drastically improve someone's life?	Williams, 1973; Singer, 1972.
Withholding reparations for slavery	Is it morally permissible for the U.S. government to refuse to pay reparations to descendants of slaves?	Morris, 1984; Boxill, 1972; Shiffrin, 2009.
Immigration restrictions	Is it morally permissible for governments to coercively prevent people from immigrating to their country?	Joshi, 2021; Hidalgo, 2021.
Long-haul vacation flights	Is it morally permissible to contribute to climate change by taking long-haul flights for fun?	Sinnott-Armstrong, 2005; Broome, 2012.
Medical testing with animals	Is it morally permissible to test medical interventions on animals to determine safety and efficacy for humans?	Cohen, 1986; Norcross, 2004.
The death penalty	Is it morally permissible to hold people responsible for wrongful behavior, even if their wrongful behavior was caused by factors outside their control (e.g., genes and culture)?	Strawson, 1994; Moody-Adams, 1994.
Euthanasia	Is it morally permissible to autonomously end one's own life to avoid suffering? Or is this a violation of human dignity?	Hardwig, 1997; Velleman, 1999.
Having more than one child	Is it morally permissible to contribute to climate change by having more than one child?	Sinnott-Armstrong, 2005; Broome, 2012.
Abortion	Do fetuses have a right not to be killed? Do mothers have a right to control what happens in and to their body?	Marquis, 1989; Moller, 2011; Thomson, 1976.
Having more than two children	Is it morally permissible to contribute to climate change by having more than two children?	Sinnott-Armstrong, 2005; Broome, 2012.
Pursuing high pay or passion rather than a career that improves the world	In choosing a career, is it morally permissible to privilege high pay or passion, even if that means foregoing a career that makes the world a better place?	Williams, 1973; Singer, 1972.

Note. Table shows moral controversies used in the course and the papers that in-class discussions were based on. The stimuli used the exact wording in the ‘Controversial Action’ column—the ‘Elaboration’ column is provided for further context concerning what the issue entails.

Participants also completed four individual difference scales in random order: Actively Open-Minded Thinking about Evidence (AOT-E; Pennycook, Cheyne, Koehler, & Fugelsang, 2020), abridged versions of the Moralized Rationality and the Importance of Rationality (MR & IR; Ståhl, Zaai, & Skitka, 2016), and the Unified Scale to Assess Individual Differences in Intuition and Deliberation (USID; Pachur & Spaar, 2015). We selected three items from the Moralized Rationality and the Importance of Rationality scales, and six items from USID, to minimize redundancy (see Supplementary Materials, S1).

Participants in PHI were additionally assigned to one of two conditions: Half of students reflected on biasing influences of intuition as part of their weekly discussions; the others reflected on course content for that week; we included six measures to investigate the effects of this intervention.

3. Results

To investigate the effect of taking moral philosophy, we contrasted PHI students with those in the control group in linear mixed-effects regressions, with random effects for items and participants. For the view measure, we were interested in the magnitude of change (vs. the direction), so we regressed the absolute differences in views across time on group assignment. For the meta-cognitive judgments, we had directional predictions, so we regressed pre- and post-semester judgments on group assignment, time, and their interaction. We tried to fit maximally complex models (with random slopes), though for most analyses, just the intercepts-only models converged, which can increase the rate of false positives (see Barr, Levy, Scheepers, & Tily, 2013; Judd, Westfall, & Kenny, 2012).

These regressions revealed significant effects of taking the moral philosophy course for all four key measures. Across controversies, taking the course resulted in more moral view change ($b = 0.26$, $t = 3.99$, $p < .001$, 95% CI = [0.13, 0.38]), greater increase in reliance on deliberation ($b = 0.58$, $t = 7.33$, $p < .001$, 95% CI = [0.43, 0.74]), greater decrease in reliance on intuition ($b = -0.23$, $t = -2.65$, $p = .008$, 95% CI = [-0.39, -0.06]), and greater increase in confidence ($b = 0.34$, $t = 4.36$, $p < .001$, 95% CI = [0.19, 0.49]), compared to control.¹ See Fig. 1 for average




Fig. 1. Average change in moral views and meta-cognitive judgments.

Note. Figure displays changes in measures from the start of the class (T1) to the end of the class (T2). Note that View displays absolute change, as there is no expected direction of change across vignettes. As all ratings occurred on 7-point scales, the full range of possible values was 0 to 6 for absolute change in view, and -6 to 6 for the remaining three variables, which reflect directional change. Regressions (Supplementary Materials S2) show all cross-group comparisons to be significant (and all moral philosophy estimates to be significantly different from 0); *** denotes significance of the time-group interaction at $p < .001$, ** at $p < .01$.

¹ Note that the last three terms correspond to interactions rather than main effects. Degrees of freedom for mixed models are not shown as they are approximated through Satterthwaite's method—see Supplementary Materials S2.

changes, Figs. 2, 3 for cross-vignette changes, and Supplementary Materials S2 for regression tables.

These results reveal that taking moral philosophy changed students' views about some moral controversies. Moreover, exploratory analyses reveal that these changes were substantial: Judgments were significantly more likely to remain unchanged from T1 to T2 for participants in the control group (relative to PHI), while judgments were significantly more likely to flip for participants in PHI (relative to control; see Fig. 4).

What drove students in PHI to change their moral views? Cross-group differences could reflect selection effects; that is, students in PHI may have distinct psychological profiles (e.g., being more open-minded). However, analyses across all items included as individual difference measures—as well as comparisons of initial moral views on controversies—reveal no systematic differences across groups. This suggests that selection effects are unlikely to explain our results (see Supplementary Materials S5).²

A more compelling explanation is that philosophy facilitates critical thinking about moral issues by exposing students to arguments and evidence—students may learn to rely more on their deliberation, or discount their intuition. In line with the intuition explanation, changes in intuition were significantly associated with absolute changes in views ($b = -0.04$, $t = -3.01$, $p = .002$, 95% CI [-0.07, -0.02]), whereas changes in deliberation ($b = 0.00$, $t = 0.20$, $p = .84$, 95% CI [-0.02, 0.04]) and confidence ($b = 0.00$, $t = -0.14$, $p = .88$, 95% CI [-0.04, 0.03]) were not (see Supplementary Materials S6). In other words, reduced reliance on intuition from T1 to T2 was the only reliable predictor of a change in moral views.

Additional analyses suggest that taking moral philosophy led some students who initially relied heavily on intuition to discount intuition almost entirely, resulting in disproportionately large changes in moral views (i.e., initial reliance on intuition moderated the relation between change in intuition and view change, $b = 0.11$, $t = 3.41$, $p < .001$, 95% CI [0.05, 0.17], see Fig. 5; Supplementary Materials S6).

Most strictly, the hypothesis that philosophy instruction led (some) students to discount intuition predicts that the relation between change in views and reliance on intuition should be unique to PHI students. An additional regression reveals this to be the case: only the relationship between intuition and view change was significantly moderated by group ($b = -0.07$, $t = -2.38$, $p = .018$, 95% CI [-0.13, -0.01]).³

Intuition may be uniquely associated with change in moral views—but in what way? Additional analyses reveal that decreased reliance on intuition was uniquely associated with *flips* of opinion—specifically for PHI students. By contrast, increasing reliance on deliberation was associated with an increase in confidence ($b = 0.22$, $t = 8.76$, $p < .001$, 95% CI [0.16, 0.28]), with a larger effect for PHI students than control ($b = 0.19$, $t = 7.44$, $p < .001$; 95% CI [0.12, 0.26]; see Supplementary Materials S10).

An important limitation is that these regressions aggregate data across controversies, and therefore cannot answer questions at the level of individual issues. More fine-grained claims are limited by our statistical power: simulations indicate an upper bound of 20% power to detect a relationship between intuition change and view change at the individual vignette level, whereas at the aggregate level, our power for

² A similar worry is that our control group's gender imbalance may confound results. However, all key results we report remain significant after controlling for gender, indicating that this is an unlikely explanation (analyses available in our repository).

³ Furthermore, changes in intuition and deliberation are weakly correlated in our data ($r = -0.20$, $t(2386) = -9.91$, $p < .001$), indicating that it is unlikely for intuition to be suppressing the effects of deliberation, or vice versa. Relatedly, it is possible that changes in intuition and deliberation interact in predicting change in views—but a fourth regression with this higher-order interaction revealed no evidence for any effect besides intuition (see Supplementary Materials S6).




Fig. 2. Change in moral views and meta-cognitive judgments for moral philosophy students.
Note. Measurement time indicates data collected before a student began the course [1] vs. at the end of the course [2]. Each column corresponds to a different moral controversy (described in Table 1), and each row shows one of the four items described in the materials section. The center of the crossbar displays the mean, error bars show bootstrapped 95% confidence intervals, and the violin plots depict smoothed densities. Asterisks show significance of Benjamini-Hochberg multiple-comparisons corrected paired *t*-tests at critical values of 0.05 (i.e., $p < .017$, indicated by *), 0.01 (i.e., $p < .004$, indicated by **), and 0.001 (i.e., $p < .0002$, indicated by ***; see Benjamini & Hochberg, 1995). We note that while the numerical difference in intuition across time went in the predicted direction for 10 of 12 vignettes, the only vignette that individually reached significance was animal testing.




Fig. 3. Change in moral views and meta-cognitive judgments for control participants.
Note. This plot is identical to Fig. 2, with data coming from the control group participants instead of those enrolled in PHI. No paired *t*-test comparisons are significant at the same thresholds.




Fig. 4. Diagram and table of view change dynamics across all issues.
Note. Table shows percentage of judgments from each group that exhibited a given pattern of persistence or change, across all vignettes. Δ shows difference across the groups, and *** denotes significance of group as a predictor of view change dynamics in linear mixed-effects models at $p < .001$, ** at $p < .01$ (see Supplementary Materials S3 for transition matrices; S4 for regression tables). Persistence is computed as giving the same answer, Flip indicates crossing over the neutral midpoint to the other side of the scale, Suspend indicates moving to the midpoint of the scale, Decide indicates moving away from the midpoint of the scale, Polarize indicates the view becoming stronger on the same side of the scale, and Moderate indicates the view becoming weaker on the same side of the scale. These categories are mutually exclusive and exhaustive.




Fig. 5. Moral view change is moderated by extent of initial reliance on intuition for PHI students.
Note. *Intuition at T1* shows judgments of reliance on intuition collected before students took the class and *Intuition at T2* shows data from the end of the class, both on a scale from “Not at all” (1) to “Entirely” (7) for the question, “Is your judgment based on intuition / emotion?” Color of the tiles show the mean absolute view change that was measured for each combination of T1 and T2 intuition judgments; numbers on plot indicate mean view change across the cells outlined by the black squares. View changes were significantly more common for students who began with high intuition ratings at T1 and ended with low ratings at T2 (right bottom corner of PHI plot).

detecting the same relationship with mixed models is 88% (95% CI [85.8, 90.0]; $\alpha = 0.05$).⁴ We return to this in the Discussion.

An additional limitation is that the reflection intervention within PHI was ineffective or underpowered: our pre-registered analyses investigating the effect of the intervention among PHI students yielded largely null results (see Supplementary Materials S0 for details and a discussion).

⁴ We thank our reviewers for suggesting these post-hoc power analyses, which are available in our repository.

4. Discussion

In this paper, we investigated the mechanisms that promote change in moral views in the context of an extensive educational intervention. Over the course of a semester, we observed substantial change in students’ moral views across key controversies, from meat eating to reparations for slavery. In line with work questioning the primacy of deliberation in moral judgment, we found that view change was predicted exclusively by decreasing reliance on intuition, whereas increased reliance on deliberation predicted increasing confidence in moral judgment (c.f., Stanley, Dougherty, Yang, Henne, & De Brigard, 2018). At the same time, our findings suggest an alternative mechanism through which reasoning and evidence can generate moral change: Calling into question the reliability of intuition and emotion.

Since antiquity, philosophers have relied on systematic doubt, reflection, and argumentation to examine the validity of their intuitive beliefs (Kraut, 2022), and this tradition survives in introductory courses. In PHI, thought experiments and evidence favoring and opposing the moral permissibility of various actions were presented in tandem. Students with intuitive beliefs were thus given reason to question the reliability of their moral judgments, whereas students with weak prior views may have been left unmoved by the balance of evidence. Our findings support this claim: the relationship between intuition and view change was driven by students whose views were based most strongly on intuition. Note that this result is unanticipated by strictly deliberative accounts, as it suggests that intuition plays a primary role in moral change. However, it is also unanticipated by strictly intuitionist accounts, which would predict the greatest change in students with the weakest reliance on intuition (c.f., Haidt, 2001). Explaining this result instead requires reconciling the two perspectives, and thus raises new questions about the interplay between intuition, deliberation, and moral judgment.

For instance, though our results show that changes in reliance on intuition play a unique role in changing views, we did not investigate which aspects of the course (e.g., thought experiments) drove this change, nor how deliberative processes may have reduced reliance on intuition (see Lombrozo, 2009). Similarly, since our sample was comprised of undergraduates in the U.S, we do not know whether particular demographic factors (e.g., age) are key to facilitating these effects (see Henrich, Heine, & Norenzayan, 2010), and we lack sufficient power for vignette-level analyses (e.g., we cannot answer whether intuition drove view change across all issues, or a subset). Finally, we note that our measures are based on self-report. Given that recent work has shown a close correspondence between self-reported moral judgment and behavior (Schwitzgebel et al., 2020) and revealed nuanced lay theories of intuition and deliberation (Oktar & Lombrozo, 2022a), our findings likely reflect consequential and sophisticated—albeit potentially biased—judgments. Taken together, these limitations highlight opportunities for future work to clarify bounds on the generalizability of our findings by investigating how different features of courses, samples, controversies, and measures moderate the effects we report.

Though deliberative aspects of the course likely drove lower reliance on intuition (and consequently view change), non-deliberative factors could also explain our results: For instance, students may have adopted the views of their peers or instructor. However, on issues with the largest change in views (meat eating and luxury spending), initial consensus dissolved over time, indicating that social contagion is unlikely to explain our findings. That the instructor kept his views private and personally disagreed with shifts on some issues speaks against the latter possibility. However, the classroom environment likely amplified the effects we report (e.g., by increasing engagement).

Our results raise additional questions: What explains cross-controversy variation in view change? Whereas judgments on some controversies shifted drastically (e.g., meat eating), judgments about others (e.g., abortion) barely budged. Though ceiling effects likely played a role (see Supplementary Fig. 1, 2), attitudes on some issues may

be seen as protected values (Sloman & Vives, 2022; Tetlock, 2003) or as common-sensical (Goodwin & Darley, 2012), shielding intuitions from counterarguments. Cultural context could also potentiate view change: For example, climate change activists have been drawing attention to the consequences of meat consumption (Broad, 2018) and the effective altruism movement has highlighted the importance of donating over luxury spending (Singer, 2009). A combination of these explanations—ceiling effects, protected values, and cultural context—likely explains much cross-controversy variance in view change.

That we found robust effects of a deliberative intervention in this study, while past psychology literature has found mixed effects, ought to be less surprising in light of the reasons that motivated this study. Moral judgments have important societal consequences (Kitcher, 2021) and psychological functions (Skitka, Hanson, Morgan, & Wisneski, 2021). Given this import, expecting the short, small, and artificial interventions typically used in moral psychology research to have predictable and large effects is optimistic (Tosh et al., 2021). A months-long intervention exposing students to strong arguments on important issues, however, may be a good starting place for evaluating the plasticity of moral views and its psychological basis.

An important implication of our results is that we may be able to promote change in moral views by providing reasons and evidence to doubt the reliability of our intuitive, emotional reactions. This conclusion is consistent with research on the unique role that beliefs about intuition play in guiding judgments and decisions across domains (Oktar & Lombrozo, 2022a; Plunkett, Buchak, & Lombrozo, 2020), and provides a hopeful counterpoint to findings on people's tendency to persist in their moral views amid mass disagreement (Oktar & Lombrozo, 2022b). Taken together, this work suggests that moral education can help sway intuitively anchored moral views—and hopefully promote moral progress—one argument at a time.

CRedit authorship contribution statement

Kerem Oktar: Conceptualization, Methodology, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **Adam Lerner:** Conceptualization, Methodology, Writing – review & editing, Investigation. **Maya Malaviya:** Investigation, Visualization, Project administration, Writing – review & editing. **Tania Lombrozo:** Conceptualization, Methodology, Investigation, Writing – review & editing, Visualization, Supervision, Funding acquisition.

Declaration of Competing Interest

We have no known conflicts of interest to disclose.

All pre-registrations, materials, analysis scripts, and data available at <https://osf.io/y5tdu/>.

We are grateful to Eric Schwitzgebel and members of the Concepts and Cognition Lab for valuable feedback on this research.

Data availability

All data and analysis scripts are available in our open access repository, at: <https://osf.io/y5tdu/>

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.cognition.2023.105434>.

References

- Bago, B., Kovacs, M., Protzko, J., Nagy, T., Kececs, Z., Palfi, B., ... Aczel, B. (2022). Situational factors shape moral judgements in the trolley dilemma in eastern, southern and Western countries in a culturally diverse sample. *Nature Human Behavior*, 1–16. <https://doi.org/10.1038/s41562-022-01319-5>

- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1), 289–300.
- Bloom, P. (2010). How do morals change? *Nature*, 464(7288), 490. <https://doi.org/10.1038/464490a>
- Boxill, B. R. (1972). The morality of reparation. *Social Theory and Practice*, 2(1), 113–123. <https://doi.org/10.5840/soctheorpract1972216>
- Broad, G. M. (2018). Effective animal advocacy: Effective altruism, the social economy, and the animal protection movement. *Agriculture and Human Values*, 35(4), 777–789.
- Broome, J. (2012). *The public and private morality of climate change*.
- Cohen, C. (1986). The case for the use of animals in biomedical research. *New England Journal of Medicine*, 315(14), 865–870. <https://doi.org/10.1056/NEJM198610023151405>
- Fletcher-Brown, D., Buono, A. F., Frederick, R., Hall, G., & Sultan, J. (2012). A longitudinal study of the effectiveness of business ethics education: Establishing the baseline. *Journal of Academic Ethics*, 10(1), 45–56. <https://doi.org/10.1007/s10805-012-9149-4>
- 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(1), 250–256. <https://doi.org/10.1016/j.jesp.2011.08.006>
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108(4), 814–834. <https://doi.org/10.1037/0033-295X.108.4.814>
- Hardwig, J. (1997). Is there a duty to die? *Hastings Center Report*, 27(2), 34–42. <https://doi.org/10.2307/3527626>
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2–3), 61–83. <https://doi.org/10.1017/S0140525X0999152X>
- Herec, J., Sykora, J., Brahmi, K., Vondracek, D., Dobesova, O., Smelik, M., ... Prochazka, J. (2022). Reflection and reasoning in moral judgment: Two preregistered replications of Paxton, Ungar, and Greene (2012). *Cognitive Science*, 46(7), Article e13168. <https://doi.org/10.1111/cogs.13168>
- Hidalgo, J. (2021). *Open borders. In Living ethics: An introduction with readings*. New York: Oxford University Press.
- Horne, Z., Powell, D., & Hummel, J. (2015). A single counterexample leads to moral belief revision. *Cognitive Science*, 39(8), 1950–1964. <https://doi.org/10.1111/cogs.12223>
- Joshi, H. (2021). For (some) immigration restrictions. In B. Fischer (Ed.), *Ethics left and right: The moral issues that divide us*. Oxford University Press.
- 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(1), 54.
- Kitcher, P. (2021). *Moral progress*. Oxford University Press.
- Kraut, R. (2022). Plato. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2022). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2022/entries/plato/>
- Lombrozo, T. (2009). The role of moral commitments in moral judgment. *Cognitive Science*, 33(2), 273–286. <https://doi.org/10.1111/j.1551-6709.2009.01013.x>
- Marquis, D. (1989). Why abortion is immoral. *The Journal of Philosophy*, 86(4), 183–202. <https://doi.org/10.2307/2026961>
- Moller, D. (2011). Abortion and moral risk. *Philosophy*, 86(3), 425–443. <https://doi.org/10.1017/S0031819111000222>
- Moody-Adams, M. M. (1994). Culture, responsibility, and affected ignorance. *Ethics*, 104(2), 291–309.
- Morris, C. W. (1984). Existential limits to the rectification of past wrongs. *American Philosophical Quarterly*, 21(2), 175–182.
- Mumford, M. D., Connelly, S., Brown, R. P., Murphy, S. T., Hill, J. H., Antes, A. L., ... Devenport, L. D. (2008). A Sensemaking approach to ethics training for scientists: Preliminary evidence of training effectiveness. *Ethics & Behavior*, 18(4), 315–339. <https://doi.org/10.1080/10508420802487815>
- Norcross, A. (2004). Puppies, pigs, and people: Eating meat and marginal cases. *Philosophical Perspectives*, 18, 229–245.
- Nucci, L. P., Narvaez, D., & Krettenauer, T. (Eds.). (2014). *Handbook of moral and character education*. New York: Routledge.
- Oktar, K., & Lombrozo, T. (2022a). Deciding to be authentic: Intuition is favored over deliberation when authenticity matters. *Cognition*, 223, Article 105021. <https://doi.org/10.1016/j.cognition.2022.105021>
- Oktar, K., & Lombrozo, T. (2022b). Mechanisms of belief persistence in the face of societal disagreement. In *Proceedings of the 44th Annual Meeting of the Cognitive Science Society*.
- Pachur, T., & Spaar, M. (2015). Domain-specific preferences for intuition and deliberation in decision making. *Modeling and Aiding Intuition in Organizational Decision Making*, 4(3), 303–311. <https://doi.org/10.1016/j.jarmac.2015.07.006>
- Parfit, D. (2011). *On what matters* (Vol. 1). Oxford University Press.
- Paxton, J. M., Ungar, L., & Greene, J. D. (2012). Reflection and reasoning in moral judgment. *Cognitive Science*, 36(1), 163–177. <https://doi.org/10.1111/j.1551-6709.2011.01210.x>
- Pennycook, G., Cheyne, J. A., Koehler, D. J., & Fugelsang, J. A. (2020). On the belief that beliefs should change according to evidence: Implications for conspiratorial, moral, paranormal, political, religious, and science beliefs. *Judgment and Decision making*, 15(4).
- Plunkett, D., Buchak, L., & Lombrozo, T. (2020). When and why people think beliefs are “debunked” by scientific explanations of their origins. *Mind & Language*, 35(1), 3–28.

- Ritter, B. A. (2006). Can business ethics be trained? A study of the ethical decision-making process in business students. *Journal of Business Ethics*, 68(2), 153–164. <https://doi.org/10.1007/s10551-006-9062-0>
- Schwitzgebel, E., Cokolet, B., & Singer, P. (2020). Do ethics classes influence student behavior? Case study: Teaching the ethics of eating meat. *Cognition*, 203, Article 104397. <https://doi.org/10.1016/j.cognition.2020.104397>
- Shiffrin, S. V. (2009). Reparations for U.S. slavery and justice over time. In M. A. Roberts, & D. T. Wasserman (Eds.), *Harming future persons: Ethics, genetics and the nonidentity problem* (pp. 333–339). Netherlands: Springer. https://doi.org/10.1007/978-1-4020-5697-0_16.
- Singer, P. (1972). Famine, affluence, and morality. *Philosophy & Public Affairs*, 1(3), 229–243.
- Singer, P. (2009). *The life you can save: Acting now to end world poverty*. New York: Random House.
- Sinnott-Armstrong, W. (2005). It's not my fault: Global warming and individual moral obligations. In W. Sinnott-Armstrong, & R. Howarth (Eds.), *Perspectives on climate change* (pp. 221–253). Elsevier.
- Skitka, L. J., Hanson, B. E., Morgan, G. S., & Wisneski, D. C. (2021). The psychology of moral conviction. *Annual Review of Psychology*, 72(1), 347–366. <https://doi.org/10.1146/annurev-psych-063020-030612>
- Sloman, S. A., & Vives, M.-L. (2022). Is political extremism supported by an illusion of understanding? *Cognition*, 225, Article 105146. <https://doi.org/10.1016/j.cognition.2022.105146>
- Ståhl, T., Zaal, M. P., & Skitka, L. J. (2016). Moralized rationality: Relying on logic and evidence in the formation and evaluation of belief can be seen as a moral issue. *PLoS One*, 11(11), Article e0166332. <https://doi.org/10.1371/journal.pone.0166332>
- Stanley, M. L., Dougherty, A. M., Yang, B. W., Henne, P., & De Brigard, F. (2018). Reasons probably won't change your mind: The role of reasons in revising moral decisions. *Journal of Experimental Psychology: General*, 147(7), 962–987. <https://doi.org/10.1037/xge0000368>
- Strawson, G. (1994). The impossibility of moral responsibility. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 75(1/2), 5–24.
- Tetlock, P. E. (2003). Thinking the unthinkable: Sacred values and taboo cognitions. *Trends in Cognitive Sciences*, 7(7), 320–324. [https://doi.org/10.1016/s1364-6613\(03\)00135-9](https://doi.org/10.1016/s1364-6613(03)00135-9)
- Thomson, J. J. (1976). A defense of abortion. In J. M. Humber, & R. F. Almeder (Eds.), *Biomedical ethics and the law* (pp. 39–54). Springer US. https://doi.org/10.1007/978-1-4684-2223-8_5.
- Tosh, C., Greengard, P., Goodrich, B., Gelman, A., Vehtari, A., & Hsu, D. (2021). The piranha problem: Large effects swimming in a small pond (arXiv:2105.13445). *arXiv*. <https://doi.org/10.48550/arXiv.2105.13445>
- Velleman, J. D. (1999). A right of self-termination? *Ethics*, 109(3), 606–628. <https://doi.org/10.1086/233924>
- Waples, E. P., Antes, A. L., Murphy, S. T., Connelly, S., & Mumford, M. D. (2009). A Meta-analytic investigation of business ethics instruction. *Journal of Business Ethics*, 87(1), 133–151. <https://doi.org/10.1007/s10551-008-9875-0>
- Watts, L. L., Medeiros, K. E., Mulhearn, T. J., Steele, L. M., Connelly, S., & Mumford, M. D. (2017). Are ethics training programs improving? A meta-analytic review of past and present ethics instruction in the sciences. *Ethics & Behavior*, 27(5), 351–384. <https://doi.org/10.1080/10508422.2016.1182025>
- Williams, B. (1973). A critique of utilitarianism. *Utilitarianism: For and Against*, 93, 97–99.