

Crowdsourcing effective educational interventions

J. Hunter Priniski (priniski@ucla.edu)

Department of Psychology, University of California, Los Angeles
Los Angeles, CA, USA

Zachary Horne (zachary.horne@asu.edu)

School of Social and Behavioral Sciences, Arizona State University
Phoenix, AZ, USA

Abstract

Creating effective educational interventions that correct people's misconceptions is difficult. This has led many researchers to conclude that people do not properly attend to new information in a way that they should. However, even if a scientifically-grounded intervention fails, it is still possible that other interventions would be effective. Yet, it is not practically feasible to systematically explore and test the entire hypothesis space of possible interventions. Here, we examined whether researchers could use online arguments to develop effective educational interventions, in effect, narrowing the intervention hypothesis space. Across two experiments ($N = 1,816$), we found that arguments crowdsourced from Reddit's Change My View were as effective or more effective at changing beliefs than interventions developed by academics and published in top-tier scientific journals. These results suggest that researchers can build on successful crowdsourced arguments to develop effective educational interventions likely to correct people's misconceptions in more naturalistic settings.

Keywords: belief change; crowdsourcing; crowd work

Introduction

It can be difficult to find common ground with people we disagree with. People's beliefs about polarizing issues are often deeply entrenched and evidence that counters these beliefs generally does not lead people to change their minds. This intransigence comes at a cost: Polarization is a growing problem in the United States (Pew Research Center, 2014) and widespread misinformation and misconceptions about, for example, climate change only exacerbate polarization, posing considerable challenges to society. What's more, even in situations when very few people hold a misinformed belief—such as believing that vaccines cause autism—the consequences can still have a widespread negative effect in society; this is evident from the recent resurgence of measles borne from parents refusing to vaccinate their children, citing fears that vaccines cause autism (Center for Disease Control, 2019).

To effectively educate the public, researchers have attempted to confront belief polarization and resistance to evidence by experimentally testing whether educational interventions can induce rational belief updating (e.g., Horne, Powell, Hummel, & Holyoak, 2015; Lai et al., 2014; Nyhan & Reifler, 2015; Nyhan, Reifler, Richey, & Freed, 2014; Turetsky & Sanderson, 2018). Ideally, people would always properly update their beliefs in accordance with the evidence. However, many interventions developed by scientists are ineffective (e.g., Nyhan et al., 2014), leading

researchers to conclude that people cannot change their beliefs about issues such as climate change, vaccination, or immigration.

There are several psychological explanations that might explain why educational interventions are often ineffective. First, people interpret evidence to confirm their previously-held beliefs (Klayman, 1995; Nickerson, 1998), and our strongly-held beliefs—such as political and moral beliefs—are deeply rooted in our views of ourselves (e.g., Strohinger & Nichols, 2014; Carney, Jost, Gosling, & Potter, 2008), and thus are particularly resistant to change (Kahan, Peters, Wittlin, Slovic, & Ouellette, 2012). Second, even when people assimilate evidence, they do so imperfectly, requiring much more evidence than seems epistemically warranted (e.g., Priniski & Horne, 2018). Even massive education campaigns seem to yield only minor changes in public opinion and behavior (e.g., Fiore et al., 1990). Together, these results have led many researchers to either conclude that meaningful belief change is, in a practical sense, infeasible or that something other than education and evidence is needed to overcome strongly-held beliefs.

However, when an educational intervention fails to change people's misconceptions, this does not entail that other educational interventions (even similar interventions) would fail as well. It is an empirical question whether an untested intervention would turn out to be efficacious. Indeed, researchers have successfully developed effective educational interventions. For instance, Lewandowsky, Gignac, & Vaughan (2013) found that making people aware of the scientific consensus surrounding climate change using icon arrays positively affected people's beliefs. More recently, researchers have found that educational interventions can change vaccine intentions (Horne et al., 2015), correct mental health misperceptions (Turetsky & Sanderson, 2017), and address implicit racial biases, though these changes may be transient (Lai et al., 2014). However, beyond combing the academic literature, researchers have little to go on in predicting whether a given untested intervention will succeed or fail. Moreover, educational interventions are rarely tested outside of the lab, which allows for the possibility that effective educational interventions developed in the lab will fail to generalize beyond tightly controlled settings (Priniski & Horne, 2018). To complicate matters further, the hypothesis space of

possible interventions is very large (read, infinite). Consequently, it is not feasible for any given lab or even a group of labs to systematically explore the entire hypothesis space of educational interventions to determine whether a possible intervention could change people's beliefs about a given topic. A methodological advance is needed to avoid a protracted search through the intervention hypothesis space.

We propose a new method for developing educational interventions: Using successful persuasive arguments culled from online discussions (for example, from the Reddit forum Change My View). We propose that developing interventions based on existing arguments that have proven to be effective in naturalistic environments provides a compelling starting place for the development of effective educational interventions.

Change My View

Change My View is a popular Reddit forum where users post their views on issues ranging from gun control to opinions about movies. Redditors posting in this community understand that others will attempt to change their view by providing arguments opposing their beliefs (see Table S1 in Supplemental Materials, found at <https://osf.io/v54ut/>). As one would expect, some arguments are more persuasive than others and thus the variance in argument quality found on the forum provides a naturalistic resource for examining the features of effective arguments.

As a naturalistic data source, Change My View has provided several insights into how belief change occurs outside of the lab. For example, Priniski and Horne (2018) found that arguments containing more statistical language and links to news or scientific articles were more likely to change other users' strongly-held beliefs—evidence can change people's minds. Other researchers have examined the logical qualities of effective arguments on the forum (e.g., use of classical modes of persuasion: ethos, logos, pathos, Hidey, Musi, Hwang, Muresan, & McKeown, 2018). Research on Change My View has extended beyond social psychology. Computer scientists have developed computational models that extract features of argumentation, such as predicting the probability an argument is effective given linguistic features (Tan, Niculae, Danescu-Niculescu-Mizil, & Lee, 2016) or machine classifying "parts" of beliefs most amenable to change (Jo et. al, 2018).

While many researchers have examined the factors that predict belief change among Change My View users, it is unknown whether effective arguments taken from this forum would be equally effective in more controlled contexts or among a population not seeking arguments opposing their beliefs. In fact, there are several reasons why belief revision may look different on Change My View than it does in the lab. These reasons pose concern for the generalizability of effective arguments found on Change My View and need to be experimentally addressed before Change My View can be recommended as a crowdsourcing platform for effective

educational interventions.

For one, people who discuss certain topics—and particularly users on Reddit's Change My View—may be more willing to change their minds and consider evidence for an opposing argument. This may not be true for the public at large, limiting the generalizability of these prior findings. Second, people engaged in a debate on a particular topic may be more motivated to deliberate on the topics they're discussing. This fact may make online communities such as Change My View an ideal population to study central rather than peripheral routes to persuasion (Petty & Cacioppo, 1986). However, it may also make online communities unrepresentative of the general population who may not be so ready to entertain evidence that is contrary to their beliefs.

Altogether, controlled laboratory research is necessary to understand if the persuasive tactics deployed online can generalize to other populations and, in turn, serve as a starting place for developing educational interventions.

Present Experiments

In the present experiments, we identified successful arguments on Change My View and performed a head-to-head comparison to interventions reported in academic psychology, public policy, political science, communications, and behavioral economics articles—adopting a methodological approach most analogous to a strategy relied on in clinical trials (e.g., Leuch, et al., 2013). Namely, we compared crowdsourced arguments to academic arguments that have been shown to be somewhat effective at changing people's beliefs (or at a minimum, exert the same task demands on participants). Performing this comparison allowed us to predict whether effective educational interventions can be culled from online communities and used as effective interventions in controlled laboratory settings.

It is worth highlighting how this experimental strategy diverges from comparing the performance of an intervention to an inactive control condition. As opposed to controlling for features of naturalistic interventions to uncover what makes them effective, the paradigm we are proposing first identifies the interventions that yield desirable consequences (e.g., a reduction in misconceptions surrounding structural racism), at which point we can subsequently uncover the mechanisms that realize these positive effects. As a consequence, academic and crowdsourced interventions will differ along many unknown dimensions (including length, the task performed, the information presented, and so on). However, we do have prior evidence (either from empirical studies or from data mined from discussion forums) that signal the efficacy of each of the interventions being compared. Ultimately, researchers aim to develop interventions that can effectively educate the public, making this dimension—efficacy—the most central on which to assess an educational intervention.

With this goal in mind, in Experiment 1 we compared the efficacy of crowdsourced and academic interventions at changing beliefs across four hotly-debated topics. Experiment 2 was an extension of Experiment 1, where we further examined whether crowdsourced arguments would be as effective as academic intervention across four new topics.

Experiment 1

Preregistration The projected sample size, predictions, and analysis scripts were preregistered through Open Science Framework. Experimental scripts, analyses, scales, and Supplemental Materials are available at <https://osf.io/v54ut/>.

Participants We recruited 916 participants through Amazon's Mechanical Turk to be 80% powered to detect a Cohen's d of .1 in a within-subjects design. Of the participants recruited, 816 passed attention checks and were included in the analysis of this study (333 men, 476 women, 4 non-binary, 4 preferred not to say; the median age of participants was 35 years old).

Interventions Participants received four separate interventions that focused on either (a) reducing racist beliefs, (b) increasing support for vaccines, (c) increasing support for gun control, and (d) reducing xenophobic attitudes directed at immigrants. Participants received two crowdsourced interventions and two academic interventions (intervention type: within-subjects) with one intervention for each topic. Therefore, we tested the efficacy of eight interventions in total. Crowdsourced interventions were copied-and-pasted comments that were awarded a "delta" in a Change My View discussion—a signification that the argument changed the view of at least one user on the forum. We selected discussion comments from Change My View as crowdsourced interventions if they met the following three criteria. First, the comment was related to a topic that psychologists have traditionally studied in the lab (e.g., climate change, gun control, xenophobia, etc.). Second, the comment had been awarded a delta. Third, the content of the comments could be developed into an intervention with little-to-no editing, content change, or manipulation. Many comments on Change My View satisfy these criteria and *could* have been empirically tested, but the aim of the present studies is to consider how several representative crowdsourced examples could be developed into effective educational interventions. (Detailed information about the interventions can be found at <https://osf.io/v54ut/>).

Pretest and Posttest Measures We examined how participants' beliefs about four controversial topics changed as a function of exposure to one of two educational interventions (crowdsourced or academic) for a given topic. Prior to completing the main portion of the study, participants answered four questions assessing their pretest beliefs about each topic. For instance, participants rated their agreement with the assertion, "Gun control in America is

ineffective at reducing overall violence and crime", which was taken from a Change My View post (in this case, a post about gun control). After responding to these four assertions, participants proceeded to the intervention and post-test portion of the experiment.

We developed four separate scales to measure people's beliefs about racism, vaccines, gun control, and xenophobia directed at immigrants. Each scale was composed of five items (with two items reverse coded). Items in a topic's posttest scale were created by rewording or expanding on a pretest assertion. For example, an item in the posttest gun control scale stated, "Societies with strict gun control have similar crime rates as societies with little to no gun control." See the Supplemental Materials for more details on pretest and posttest measures.

Procedure The experiment proceeded as follows: First, participants rated their agreement with items measuring their pretest beliefs towards all four topics. Next, participants were randomly assigned either an academic or a crowdsourced intervention for a given topic. After completing this intervention (e.g., after reading information about gun control), participants responded to that topic's posttest scale. After completing the posttest scale for a given topic, participants advanced to a new topic and the procedure was reiterated until they finished reading and responding to questions about all four topics. The ordering and exposure to a given intervention type was counterbalanced and randomized.

Results and Discussion

Analytic Approach To test our hypotheses, we performed Bayesian mixed effects modeling using the R package `brms` (Burkner, 2018). We set regularizing priors for all population-level effects in our models, which we detail below. These priors are recommended because they provide conservative effect size estimates and reduce the likelihood of overfitting (Gelman, Lee, & Guo, 2015; McElreath, 2016). Following the recommendations of Liddell & Kruschke (2018), Likert data were modeled with a cumulative probability distribution. The cumulative distribution is recommended for Likert scale data because it assumes that ordered responses represent a continuous latent construct.

We tested our hypothesis by fitting an ordinal mixed-effects model predicting posttest beliefs based on the interaction between condition (Reference = Academic condition) and topic (Reference = Guns). This model controlled for participants' responses to the pretest statement, which we treated as a monotonic effect. This model included group-level effects of Subject and Topic and allowed for heterogeneity in the slopes of the effects of Condition and Topic on participants' responses. Our model is specified below in `brms` syntax (Bürkner, 2018):

```
Response ~ Condition*Topic + mo(PreTest)
+ (1 + Topic + Condition | Subject)
```




Figure 1: Posttest responses for each intervention tested in Experiment 1 (1 = Strongly disagree; 7 = Strongly agree). Relative effectiveness of a crowdsourced intervention can be seen by comparing the leftward shift of responses across interventions for a topic. Figures S5 through S8 in the Supplemental Materials show posttest responses grouped by pretest response for each intervention tested in tested in Experiment 1.

Bayesian analyses formulate model parameters as probability distributions wherein the posterior distribution for a parameter θ is computed via the prior and likelihood of θ . To model the joint probability distribution of participants' responses, we specified priors over the possible effects each parameter could have on our response variable:

$$\begin{aligned}
 \beta_{Intercept[1]} &\sim \mathcal{N}(2.19, 1) \\
 \beta_{Intercept[2]} &\sim \mathcal{N}(2.94, 1) \\
 \beta_{Intercept[3]} &\sim \mathcal{N}(3.17, 1) \\
 \beta_{Intercept[4]} &\sim \mathcal{N}(3.47, 1) \\
 \beta_{Intercept[5]} &\sim \mathcal{N}(3.89, 1) \\
 \beta_{Intercept[6]} &\sim \mathcal{N}(4.59, 1) \\
 \beta_{Condition} &\sim \mathcal{N}(0, .5) \\
 \beta_{Pretest\ Beliefs} &\sim \mathcal{N}(4, 2) \\
 \beta_{Topics} &\sim \mathcal{N}(0, 3) \\
 \beta_{Topic \times Condition\ Interactions} &\sim \mathcal{N}(0, .5)
 \end{aligned}$$

$\Omega_k \sim LKJ(1)$ where Ω_k is a correlation matrix of group-level parameter
Group-level parameters $\sim \mathcal{N}(1, 2)$

These analyses revealed that the crowdsourced interventions countering racist ($b = -0.58$, 95% CI [-.80, -.37]) and anti-immigrant beliefs ($b = -.40$, 95% CI [-.60, -.18]) were credibly more effective than an academic intervention; interventions on vaccines and gun control were equally effective (see Figure 1). These results suggest that there are arguments being developed in online communities that are comparably effective to interventions behavioral scientists have developed. And considering crowdsourced arguments have the additional virtue of being shown to be effective in a naturalistic setting free from task demands, this may give additional motivation for beginning development of educational interventions on the basis of crowdsourced arguments.

However, given that the present design lacks a completely neutral control condition, it is important to be clear on what these results do not show. First, these results do not demonstrate the true magnitude of the effect of a given intervention. Second, there is a large amount of variance in intervention quality and effectiveness for any intervention type, and there is no reason to think that all crowdsourced arguments will always be as effective or more effective than academic interventions. Rather, one should interpret the results of Experiment 1 as suggesting that crowdsourced arguments can provide a starting place for developing educational interventions and doing so has the additional virtue of giving us a priori reason to think they will generalize to comparatively more naturalistic settings.

Experiment 2

Experiment 2 was a preregistered extension of Experiment 1. The registration for this project can be found at <https://osf.io/v54ut/>. This experiment followed an identical procedure but tested the efficacy of academic and crowdsourced interventions on four new topics: (a) reducing sexist beliefs, (b) reducing transphobic beliefs, (c) reducing denial in the negative effects of climate change, and (d) reducing favor for capital punishment.

Participants We recruited 900 participants through Amazon's Mechanical Turk to be 80% powered to detect a Cohen's d of .1 in a within-subjects design. Of the participants recruited, 745 passed attention checks and were included in the analysis of this study (325 men, 416 women, 3 non-binary, 1 preferred not to say; the median age of participants was 33 years old).

Results and Discussion

Like Experiment 1, we predicted that crowdsourced interventions would be as effective or more effective than academic interventions for the four new topics. We fit the same ordinal regression model with the same priors as




Figure 2: Posttest responses for each intervention tested in Experiment 2 (1 = Strongly disagree; 7 = Strongly agree). Relative effectiveness of a crowdsourced intervention can be seen by comparing the leftward shift of responses across interventions for a topic. Figures S9 through S12 in the Supplemental Materials show posttest responses grouped by pretest response for each intervention tested in tested in Experiment 2.

Experiment 1.

In Experiment 2, we found that crowdsourced interventions were equally effective as academic interventions across three topics; the academic intervention aimed at shifting people’s beliefs about climate change was more effective than the crowdsourced intervention, $b = 0.24$, 95% CI [.00, .40].

Discussion

People’s beliefs about topics like science and morality are stubbornly resistant to new information. Developing educational interventions to correct these beliefs is a difficult task that often results in fruitless outcomes. It is also often unknown whether an intervention that manages to successfully shift beliefs in the lab will be similarly effective in a more naturalistic setting. The present studies suggest

that researchers can use crowdsourced arguments to better predict and develop effective educational interventions. Furthermore, crowdsourcing effective arguments can impact the study of belief revision directly by elucidating which types of information are most effective at changing strongly held beliefs: a topic of interest to many researchers studying higher-level cognitive processes. In two experiments, we tested whether arguments crowdsourced from the Reddit forum Change My View could be used to such an end. In Experiments 1 and 2, we compared arguments crowdsourced from Change My View to interventions taken from academic research in psychology, communications, political science, behavioral economics, and public policy. In Experiment 1, we found that across four topics, crowdsourced arguments were as effective or more effective at changing beliefs compared to previously published or tested educational interventions developed by academics. Experiment 2 followed the same procedure, finding that crowdsourced interventions were as effective at changing beliefs in three of four topics. In only one case did an academic intervention perform better at correcting scientific misconceptions than a crowdsourced intervention.

In light of these results, we propose that arguments mined from online communities can be used to develop educational interventions. How might this process work? Consider the results in Experiment 2: We observed that an academic intervention containing an icon array (Lewandowski, et al., 2013) was more persuasive than a similar crowdsourced intervention that did not contain data visualization. This finding is consistent with a large body of research demonstrating that data visualizations can effectively communicate complex information (e.g., Fernandes, Walls, Munson, Hullman, & Kay, 2018). In future research, we propose that researchers could begin to develop an educational intervention by first turning to crowdsourced interventions that appear effective and then extending them based on well-established theoretical considerations. For instance, we found that a crowdsourced intervention about the repercussions of structural racism was much more effective than an academic intervention aimed at shifting people’s implicit racial biases (Lai et al., 2014). One possibility, then, is that we could further improve the efficacy of this crowdsourced intervention by augmenting it with compelling visualizations. In this way, researchers would be able to develop interventions that have the twin virtues of demonstrating prior success in naturalistic environments and having strong empirical support from controlled laboratory studies.

However, the present experiments have some clear limitations. By design, both experiments lacked a true control condition, leaving an important question unanswered: Exactly how effective are these interventions at changing beliefs? The present studies compared the *relative* effectiveness of crowdsourced interventions to academic interventions, and didn’t demonstrate how effective they are

with respect to a neutral control condition. Future work should compare interventions to a true control condition in order to make explicit how effective a given intervention is at changing beliefs.

Change My View is also not the only place researchers could crowdsource effective arguments; a web application could also assist in mining, for example, Facebook and Twitter for effective arguments. The tool we are proposing could take queries (e.g., topics for an intervention) and return effective arguments filtered by the searched terms. Such a system could allow researchers to not only crowdsource educational interventions more effectively, but also gain an understanding of how arguments are communicated and received among members of online communities.

A cursory look on Reddit, Twitter, and Facebook demonstrates that people naturally engage in (sometimes) persuasive argumentation. Here, we proposed that psychologists can mine this information to efficiently create educational interventions that are more likely to persuade people than the methods researchers currently use—crowdsourced interventions have the advantage of being vetted, so to speak, in naturalistic contexts. Two experiments provide support for this proposal. We observed that crowdsourced arguments were more effective or often as effective as academic interventions aimed at correcting misconceptions about several societally important topics.

References

- Bürkner, P. C. (2018). Advanced Bayesian Multilevel Modeling with the R Package brms. *The R Journal*, 10, 395–411.
- Carney, D. R., Jost, J. T., Gosling, S. D., & Potter, J. (2008). The secret lives of liberals and conservatives: Personality profiles, interaction styles, and the things they leave behind. *Political Psychology*, 29, 807–840.
- Center for Disease Control. (2019). Measles cases in the u.s. are highest since measles was eliminated in 2000. Retrieved from cdc.gov.
- Fernandes, M., Walls, L., Munson, S., Hullman, J., & Kay, M. (2018). Uncertainty Displays Using Quantile Dotplots or CDFs Improve Transit Decision-Making. In R. Mandryk & M. Hancock (Eds.), *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. New York, NY: Association for Computing Machinery.
- Fiore, M. C., Novotny, T. E., Pierce, J. P., Giovino, G. A., Hatzianreou, E. J., Newcomb, P. A., ... Davis, R. M. (1990). Methods Used to Quit Smoking in the United States: Do Cessation Programs Help? *JAMA*, 263, 2760–2765.
- Flynn, D. J., Nyhan, B., & Reifler, J. (2017). The nature and origins of misperceptions: Understanding false and unsupported beliefs about politics. *Advances in Political Psychology*, 38, 127–150.
- Gelman, A., Lee, D., & Guo, J. (2015). Stan: A Probabilistic Programming Language for Bayesian Inference and Optimization. *Journal of Educational and Behavioral Statistics*, 40, 530–543.
- Horne, Z., Powell, D., Hummel, J. E., & Holyoak, K. J. (2015). Countering antivaccination attitudes. *Proceedings of the National Academy of Sciences*, 112, 10321–10324.
- Jo, Y., Poddar, S., Jeon, B., Shen, Q., Rose, C., & Neubig, G. (2018). Attentive Interaction Model: Modeling Changes in View in Argumentation. In M. Walker (Ed.), *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 103 – 116). Stroudsburg, PA: Association for Computational Linguistics.
- Kahan, D. M., Peters, E., Wittlin, M., Slovic, P., & Ouellette, L. L. (2012). The polarizing impact of science literacy and numeracy on perceived climate change risks. *Nature Climate Change*, 2, 732.
- Klayman, J. (1995). Varieties of confirmation bias. *Psychology of Learning and Motivation – Advances in Research and Theory*, 32, 285–418.
- Lai, C. K., Marini, M., Lehr, S. A., Cerruti, C., Shin, J. E., Joy-Gaba, J. A., ... Nosek, B. A. (2014). Reducing Implicit Racial Preferences: I. A Comparative Investigation of 17 Interventions. *Journal of Experimental Psychology: General*, 143, 1765–1785.
- Leucht, S., Cipriani, A., Spineli, L., Mavridis, D., Orey, D., Richter, F., ... Davis, J. M. (2013). Comparative efficacy and tolerability of 15 antipsychotic drugs in schizophrenia: a multiple-treatments meta-analysis. *The Lancet*, 382, 940.
- Lewandowsky, S., Ginac, G. E., & Vaughan, S. (2013). The pivotal role of perceived scientific consensus in acceptance of science. *Nature Climate Change*, 3, 399 – 404.
- Liddell, T. M., & Kruschke, J. (2017). Analyzing ordinal data with metric models: What could possibly go wrong? *Journal of Experimental Social Psychology*, 79, 328–348.
- Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37, 2098 – 2109.
- McElreath, R. (2016). *Statistical rethinking* (1st ed.). Boca Raton : CRC Press/Taylor & Francis Group.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2, 175 – 220.
- Nyhan, B., & Reifler, J. (2015). Does correcting myths about the flu vaccine work? an experimental evaluation of the effects of corrective information. *Vaccine*, 33, 459 – 464.
- Nyhan, B., Reifler, J., Richey, S., & Freed, G. L. (2014). Effective messages in vaccine promotion: A randomized trial. *Pediatrics*, 133, 835 – 842.
- Pew Research Center. (2014). Political Polarization in the American Public. Retrieved from people-press.org.
- Priniski, J. H., & Horne, Z. (2018). Attitude Change on Reddit's Change My View. In T. T. Rogers, M. Rau,

- X. Zhu, & C. W. Kalish (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 2276 – 2281). Austin, TX: Cognitive Science Society.
- Strohminger, N., & Nichols, S. (2014). The essential moral self. *Cognition*, *131*, 159 – 171.
- Tan, C., Niculae, V., Danescu-Niculescu-Mizil, C., & Lee, L. (2016). Winning arguments: Interaction dynamics and persuasion strategies in good-faith online discussions. In J. Bourdeau, J. A. Hendler, & R. N. Nkambou (Eds.), *Proceedings of the 25th International Conference on World Wide Web* (pp. 613 – 624). New York, NY: Association for Computing Machinery.
- Turetsky, K. M., & Sanderson, C. A. (2017). Comparing educational interventions: Correcting misperceived norms improves college students' mental health attitudes. *Journal of Applied Social Psychology*, *48*, 46 – 55.