

Attitude Change on Reddit's Change My View

J. Hunter Priniski (john.priniski@asu.edu)

Arizona State University,
School of Social and Behavioral Sciences, 13591 N. 47th Ave
Phoenix, AZ 85306 USA

Zachary Horne (zachary.horne@asu.edu)

Arizona State University,
School of Social and Behavioral Sciences, 13591 N. 47th Ave
Phoenix, AZ 85306 USA

Abstract

People generally ignore evidence that is contrary to their beliefs (Nickerson, 1998). To examine the factors that promotes attitude change with a new perspective, this study examined how people change their beliefs on a range of topics from gender identity to gun control on the Reddit forum *Change My View*. Specifically, we examined how people on Change My View cite evidence to change other people's minds. As prior work suggests, we find that people are not easily convinced to change their beliefs about social and moral issues, and this occurs even though people cite considerably more evidence while discussing these issues. However, our data provides one source of optimism: We found that the amount of evidence provided in a discussion predicts attitude change, suggesting that while attitude change is hard-won, providing facts and evidence may nonetheless be an effective persuasive tactic.

Keywords: attitude change, Reddit, Change My View, naturalistic data

Introduction

When asked about his views on the current state of political polarization in the United States, former President Obama stated that "one of the biggest challenges...to our democracy is the degree to which we don't share a common baseline of facts" (Abramson, 2018). As President Obama noted, it often seems like people differ not only in what they believe to be true or false, but also what evidence deserves to be the basis of our beliefs. Evidence, facts, scientific research and the like, as President Obama discussed, seem to play little role or are considered irrelevant in how people form and revise their beliefs.

As one might expect, empirical research on attitude change echoes our everyday sense of how difficult it is to find common ground with people we disagree with. Prior work suggests that, for instance, individuals interpret evidence to confirm their beliefs (Klayman, 1995; Nickerson, 1998), and the foundations of beliefs about issues most central to our identities, such as our moral or political beliefs, are deeply rooted in our views about who we are (e.g., Strohminger & Nichols, 2014, 2015; Carney,

Jost, Gosling, & Potter, 2008) and thus are highly resistant to being changed (Kahan, Peters, Wittlin, Slovic, & Ouellette, 2012).

Given that belief polarization and resistance to evidence may pose significant challenges to society (Flynn, Nyhan, & Reifler, 2017), understanding the factors that promote attitude change is necessary to making social progress; rational attitude change would ideally lead people to seeing otherwise polarizing issues in a similar light and thus they may be more likely to find common ground. To understand the factors that promote attitude change, researchers have conducted controlled laboratory studies on select subsets of topics about which people tend to be most polarized (e.g., political beliefs, Hameiri, Porat, Bar-Tal, Bieler, & Halperin, 2014; beliefs about race, Stewart & Payne, 2008). Typically, this research is conducted by providing participants with information for or against their beliefs in relatively contrived environments (e.g., Nyhan, Reifler, Richey, & Freed, 2014; Nyhan & Reifler, 2015; Horne, Powell, Hummel, & Holyoak, 2015; Lewandowsky, Oberauer & Gignac, 2013). Although this method yields tight control, it also makes it hard to determine how efficacious these interventions would be on a wider-scale and in more naturalistic environments. Furthermore, these studies focus on changing beliefs on individual, highly specific, topics (e.g., political beliefs, vaccines, climate change) making it difficult to determine whether the factors that promote attitude change in one context will generalize to other domains or issues.

To understand the mechanisms underlying attitude change, we took a different tack: We studied the factors that promote attitude change by examining how people change their minds about a variety of social and moral topics in a (comparatively) more naturalistic setting. Specifically, we examined the factors that promote attitude change using data from a popular online community: Reddit's *Change My View*. This forum is a popular Reddit forum where users post their stance on issues ranging from gun control to opinions about movies with the understanding that others will attempt to change their view by providing arguments opposing their perspective (see Table 1 below).

Table 1. *Example discussion topics and argument responses on Change My View.*

Discussion topic	Example response to original post
Unpaid internships should be illegal	It's sad that unpaid internships have the effect of freezing out talented people who can't afford a few months' living expenses without generating income. However, they play an important role in any developed economy...
There is no moral justification for eating meat in a first world country	It really depends on what your morals are. If one simply does not see it as being immoral to kill an animal for its meat, then killing the animal does not conflict with that individual's morality. Unless we were to believe in some objective sense of morality, in which case a lot of animals would be considered immoral for killing other animals to eat...

As one would expect, some user-generated arguments are more persuasive than others and thus provide a naturalistic dataset for examining the factors that predict attitude change outside of the lab. However, rather than attempting to extract nuanced but not easily generalizable linguistic properties of convincing arguments in online communities (Hidey, Musi, Hwang, Muresan, & McKeown, 2017; Jo, Poddar, Jeon, Shen, Rose, & Neubig, 2018; Musi, 2017; Tan, Niculae, Danescu-Niculescu-Mizil, & Lee, 2016), our work looked at global factors of attitude change across entire discussion threads taking place online. We did this with the hope that by identifying these global factors, behavioral scientists and policymakers could incorporate the lessons of attitude change occurring in the wild into educational interventions developed in the lab.

We began to study this rich source of data by investigating how attitude change varies between posts focused on “sociomoral” and “non-sociomoral” topics. Sociomoral posts, as we are defining them, relate to social and moral issues; the most common sociomoral posts in this forum concern politics, questions of gender identity, and current events like recent elections. In contrast, posts that are not sociomoral are a grab bag of other topics including humor and debates about movies and fiction. We sought to address three questions about sociomoral attitude change, which using this sort of naturalistic dataset can uniquely address: First, as we would intuitively expect, do people change their minds less often about sociomoral issues, in general, compared to non-sociomoral issues? Second, how do the contents of arguments differ, or do they differ at all, for these two types of discussions? Finally, regardless of domain, do facts, evidence, and data promote attitude change in online forums? By answering these questions, we sought to understand the overarching factors

that promote attitude change in real life settings about a variety of topics in a way that cannot be easily studied in the lab.

Methods

Our study had five procedural components: 1) Preregistration, 2) Collecting users’ submissions on the website Change My View, 3) classifying posts as sociomoral or not 4) categorizing arguments that successfully change someone’s mind, and 5) quantifying how much evidence was provided in a given discussion thread.

Preregistration

We preregistered this project’s procedure and our hypothesis concerning attitude change in sociomoral posts on Open Science Framework. The registration for this study can be found at the following link osf.io/jdxa8.

Submission Collection

We developed a Python script that collected 500 *top* Change My View posts using Version 5.3.0 of the Python Reddit API Wrapper (PRAW) (2017). *Top* posts are rich and mature discussions with many reply threads and participating users. Analyzing *top* posts allows us to consider well-developed discussions in their entirety as opposed to “young” discussions that have few comments.

Submission Classification

We coded a post as *sociomoral* if it concerned political, moral, or social issues. Two posts coded as sociomoral in our dataset were “U.S. military spending is unnecessarily large” and “Donald Trump has drastically changed the political landscape”. Alternatively, posts that were coded as *non-sociomoral* sometimes involved fictional components or intended to be humorous. Two examples of posts coded as non-sociomoral in our dataset were “Thank You Cards are a waste of time and money” and “Luigi is the superior Mario Brother”. All 500 posts were coded by J. Priniski, and then a second hypothesis blind coder recoded 25% of the posts ($N = 125$), agreeing on 88.8% of the original codings.

Measuring Attitude Change

We sought to examine the factors that promoted attitude change on Change My View. To this end, we developed a way to flag posts that changed people’s minds. On Change My View, there is a protocol—namely, *delta awarding*—which serves as a proxy of attitude change. Both the original poster and others can award comments a *delta* if they even partially change their mind about an issue. Delta awarding occurs when a user signifies that an argument has changed their mind. A delta can be awarded by replying to

a comment with one of the following delta strings: “Δ” and “!delta”.

To find Delta Awarded Comments (DACs), we traversed the discussion tree returned by the Reddit API using breadth-first search and string matching each comment for a delta signification. When a delta string is encountered, we moved upwards through the node’s ancestors until the root of the thread is found. This allowed us to distill the thread of conversation that lead someone to change their mind. In some cases, there is a back-and-forth between, for instance, two users until one user is finally convinced of the argument. In these cases, the thread is multiple replies in length. In most cases, however, the thread is only a single reply deep. All delta threads gathered from the 500 top posts and the code that collected them can be found at osf.io/yvunj.

Measuring Evidence Use

In addition to collecting DACs, we also examined how Change My View users incorporate evidence in their replies. Since we are using naturalistic data, we were required to some extent to make some inferential leaps in just what constitutes *evidence* in Reddit forums. We calculated evidence use by considering two measures: (1) the number of hyperlinks that cite external websites and documents and (2) a discussant’s use of “statistical language.” To collect and count hyperlinks, we searched the markup text returned by the Reddit API for words containing typical website and document identifiers, such as: ‘http://’, ‘www.’, ‘.pdf’, ‘.com’, etc. The complete list of identifiers can be found on the project’s Github: github.com/jpriniski/CMV. We calculated statistical language by string matching words in discussion threads with statistical terms and symbols, such as: ‘data’, ‘%’, ‘stats’, and so on. The code that completes this task can be found on the Github linked above.

Results

Analytic Strategy

Rather than performing null hypothesis significance testing, we performed Bayesian modeling using the programming language Stan in the R package brms. We specified priors to guide estimation of the data but these priors did not predetermine the results of any analysis. All the analyses reported herein are robust to different prior choices.

Our first question concerned the ways in which discussion of sociomoral issues differ compared to discussions of non-sociomoral issues. We examined how the rate of participation differs in discussions that concerned sociomoral and non-sociomoral topics. Sociomoral issues, by definition, are related to issues relevant to society at large and thus are more likely to be of interest to many people. We measured interest and

participation by predicting the total number of comments in a discussion thread on the basis of discussion topic (i.e., sociomoral or non-sociomoral). As one might expect, we found that there was considerably more interest in sociomoral discussions compared to non-sociomoral discussions (see Table 1).

Table 2. *Poisson regression predicting the amount of comments in a discussion based on topic type.*

Effect	Estimate	95% CI	
		Lower	Upper
Intercept	5.59	5.58	5.60
Sociomoral	0.26	0.25	0.28

Note. Non-sociomoral posts were the reference group

We then examined whether sociomoral posts prompt people to cite more evidence to support their beliefs than non-sociomoral posts. We calculated the number of comments that contained links, the total amount of links in a discussion, and the total amount of “statistically-oriented language” used in the discussion. These analyses indicated that evidence is more frequently provided in people’s debates about sociomoral topics than non-sociomoral topics, see Figure 1 and Table 2.

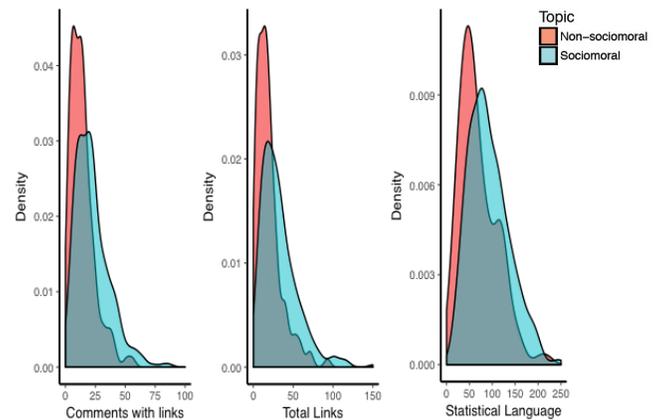


Figure 1. The number of comments using a link, the total amount of links, and the total use of statistical language for sociomoral and non-sociomoral discussions.

Table 3. *A multivariate negative binomial model predicting the amount of comments with links, the total amount of links in a discussion, and the amount of statistical language based on whether the thread concerned a sociomoral issue or not.*

Effect	Estimate	95% CI	
		Lower	Upper
Intercept	2.76	2.66	2.85
Comments using links	0.19	0.08	0.30
Intercept	3.13	3.02	3.24
Total links	0.25	0.11	0.38

Intercept	4.35	4.29	4.41
Stat. Lang.	.030	-0.04	0.11

Note. Non-sociomoral posts were the reference group

However, even though users cited considerably more evidence to advance their arguments, attitude change in the sociomoral domain is as common it is in the non-sociomoral domain – that is, more evidence yielded equivalent amounts of attitude change. As Figure 2 shows, we found that total delta awarding and delta Awarded Comments (i.e., our measures of attitude change) occur at similar rates in sociomoral and non-sociomoral threads (see Table 3).

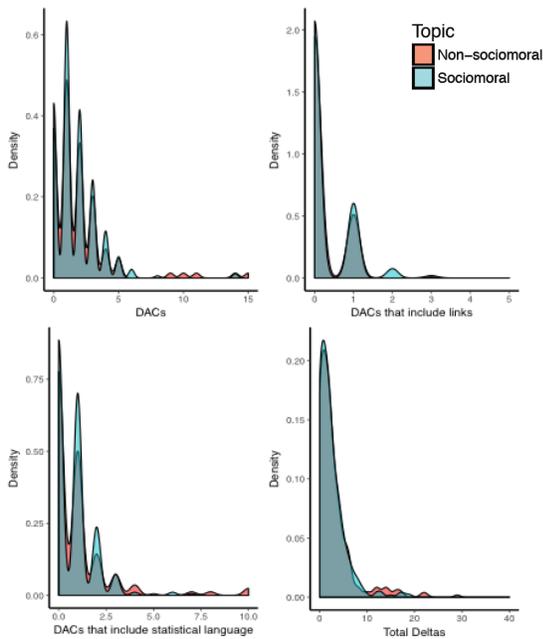


Figure 2. Density plots of how the number of Delta Awarded Comments (DACs), DACs including links, DACs using statistical language, and the total number of deltas awarded differ across sociomoral and non-sociomoral discussion topics.

Table 4. A multivariate Poisson regression predicting the number of Delta Awarded Comments (DACs), the number of DACs with links, the number of DACs that include statistical language, and the total amount of deltas on the basis of topic type.

Effect	Estimate	95% CI	
		Lower	Upper
Intercept	0.65	0.51	0.77
DACs	-0.13	-0.28	0.02
Intercept	-1.29	-1.63	-0.96
DACs including links	0.00	-0.39	0.41
Intercept	0.00	-0.18	0.17

DACs including Stat. Lang.	-0.22	-0.43	-0.01
Intercept	1.09	0.89	1.30
Deltas	-0.13	-0.37	0.11

Note. Non-sociomoral posts were the reference group.

Table 5. Attitude change predicted by the amount of evidence cited in a discussion.

Effect	Estimate	95% CI	
		Lower	Upper
<i>Effects of comments with links and total links on number of deltas</i>			
Intercept	0.55	0.31	0.79
Comments with links	0.67	0.48	0.87
Sociomoral	-0.23	-0.45	0.00
Comments	-0.12	-0.24	0.01
Intercept	0.82	0.59	1.05
Total links	0.39	0.22	0.57
Sociomoral	-0.18	-0.41	0.05
Comments	-0.01	-0.13	0.10
<i>Effects of comments with links and total links on number of DACs</i>			
Intercept	0.38	0.21	0.55
Comments with links	0.35	0.21	0.48
Sociomoral	-0.18	-0.33	-0.03
Comments	-0.12	-0.20	-0.03
Intercept	0.56	0.41	0.70
Total links	0.13	0.03	0.23
Sociomoral	-0.15	-0.30	0.01
Comments	-0.04	-0.12	0.04
<i>Effects of statistical language on number of deltas</i>			
Intercept	1.09	0.89	1.29
Stat. Lang.	0.20	0.05	0.34
Sociomoral	-0.13	-0.37	0.10
Comments	-0.02	-0.17	-0.13
<i>Effects of statistical language on number of DACs</i>			
Intercept	0.64	0.51	0.77
Stat. Lang.	0.06	-0.03	0.16
Sociomoral	-0.13	-0.28	0.03
Comments	-0.04	-0.14	0.06

Note. We distinguish between Delta Awarded Comments (DACs) and deltas because they could be distinct measures of attitude change. For example, very few comments could be awarded several deltas (one comment could receive 10 deltas). Alternatively, several comments could each be awarded just a few

deltas (three comments could each receive two deltas). For this reason, we thought it was important to distinguish these indicators of attitude change.

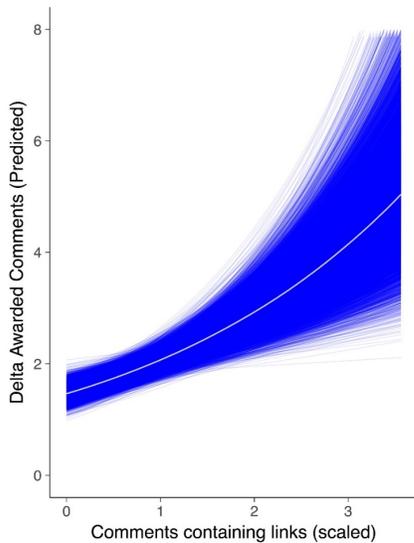


Figure 3. A marginal effects spaghetti plot predicting number of Delta Awarded Comments based on the number of comments that contained links to external sources of information. The blue region of the figure represents 95% Credibility Intervals.

These results suggest that even though commenters on Change My View participate in a forum dedicated to challenging one’s own views, the members of this community are not particularly likely to change their minds. In this way, users on Change My View are perhaps more representative of how people “in the wild” change their minds than we might initially expect.

It may appear that our results are compelling evidence for the pessimistic view that people cannot agree on what facts, if any, are even relevant to debates about, for example, politics, morality, or gender. We found that, even though substantially more evidence is cited in sociomoral discussions, attitude change is no more common in these discussions than non-sociomoral discussions.

However, our data provides some reason for optimism: Consistent with prior behavioral research (e.g., Baesler & Burgoon, 1994), we observed in a highly naturalistic dataset and across several measures of “evidence”, citing sources and referencing data was positively related to attitude change (see Figure 3 and Figure 4). When a thread contained, for instance, more citations, links to external sources, or statistical language, it positively predicted attitude change. Furthermore, we found that this effect did not depend on the discussion being sociomoral in nature (see Table 4 below; more details can be found at osf.io/s3rny). This result provides “real world” evidence that when people are motivated to attend to information relevant to their beliefs, citing sources, providing data, and so forth can be an efficacious tactic for changing people’s

attitudes (see Petty & Cacioppo, 1986, for prior laboratory based studies suggesting this same conclusion).

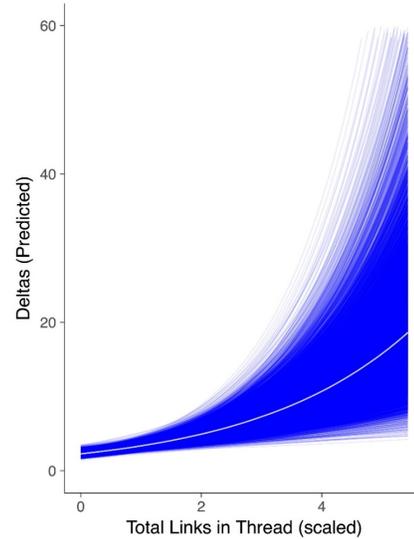


Figure 4. A marginal effects spaghetti plot predicting the total number of deltas awarded in a discussion thread based on the total number of links to external sources of information. The blue region of the figure represents 95% Credibility Intervals.

Discussion

Here, we examined the factors that promote attitude change in hotly debated topics, using a naturalistic dataset by studying attitude change in over 100,000 comments in 500 discussion threads on Reddit’s Change My View. This study revealed that even though users cite considerably more evidence while discussing sociomoral topics, they appear equally likely to revise their beliefs compared to topics that do not concern sociomoral issues. However, we also found that regardless of discussion type, providing evidence for a claim (for example, in the form of links to external articles) positively impacted people’s willingness to change their minds. Thus, our work may suggest that while attitude change is hard-won, providing facts, statistics, and citations for one’s arguments can convince people to change their minds.

Limitations and Future Directions. One concern with using Reddit’s Change My View to understand how attitudes change “in the wild” is that the people who participate in this Reddit forum may be particularly willing to change their minds and to consider statistical evidence for an argument. This may not be true in the case of the general population (e.g., Nickerson, 1998), which may limit the generalizability of our findings. Several facts speak against this concern, however. First, we observed that, in a given thread, approximately 400 comments would yield only two deltas – only two people change their mind in a thread containing numerous arguments, citations, and

statistics. This is exactly the kind of proportion we would expect to observe in the general population where everyday experience tells us that attitude change rarely occurs, if ever. Second, we also found that, as we would expect about the general population, it is harder to change people's minds over sociomoral issues compared to non-sociomoral issues. So, in contrast to this initial hunch about the users of Change My View, we suspect that our results here would be more representative of how attitude change occurs beyond the artificial conditions imposed by the laboratory.

However, there are unquestionably several limitations of this naturalistic dataset that must be acknowledge. First, it may be that Redditors are unwilling to award a delta even when they have experienced attitude change, a limitation that future research may be able to address by surveying Redditors. Second, it must be noted that members of this community are *motivated* to deliberate on things discussed in the threads on Change My View. This quality of the forum users makes it an ideal population to study central rather than peripheral routes to persuasion (Petty & Cacioppo, 1986) but may be unrepresentative of the general population. For this reason, more research is necessary to understand the extent to which the persuasive tactics deployed by users on Change My View would generalize to populations who are not so motivated to consider the facts relevant to their beliefs.

References

- Abramson, A. (2018, January 12). 'We don't share a common baseline of facts.' Barack Obama reflects on divisiveness in politics. *Time*. Retrieved from <https://www.time.com>
- Baesler, E. J., & Burgoon, J. K. (1994). The temporal effects of story and statistical evidence on belief change. *Communication Research*, 21(5), 582-602.
- 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(6), 807-840.
- 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., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2014). *Bayesian Data Analysis* (Vol. 2). London, England: Chapman & Hall/CRC.
- Hameiri, B., Porat, R., Bar-Tal, D., Bieler, A., & Halperin, E. (2014). Paradoxical thinking as a new avenue of intervention to promote peace. *Proceedings of National Academy of Sciences*, 111(30), 10995-11001.
- Hidey, C., & McKeown, K. (2018). Persuasive Influence Detection: The Role of Argument Sequencing. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*. New Orleans, Louisiana.
- Horne, Z., Powell, D., Hummel, J.E., & Holyoak, K.J. (2015). Countering antivaccination attitudes. *Proceedings of the National Academy of Sciences*, 112(33), 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 *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. New Orleans, Louisiana.
- 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(10), 732.
- Klayman, J. (1995). Varieties of confirmation bias. *Psychology of Learning and Motivation – Advances in Research and Theory*, 32(C), 385–418.
- Lewandowsky, S., Oberauer, K., & Gignac, G.E. (2013). NASA faked the moon landing – Therefore, (climate) science is a hoax: An anatomy of the motivated rejection of science. *Psychological Science*, 24(5), 622-633.
- 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(11), 2098–2109.
- Musi, E. (2017). How did you change my view? A corpus-based study of concessions' argumentative role. *Discourse Studies*, 20(2), 270-288.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(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(3), 459-464.
- Nyhan, B., Reifler, J., Richey, S., & Freed, G.L. (2014). Effective messages in vaccine promotion: A randomized trial. *Pediatrics*, 133(4), 835-842.
- Petty, R.E., & Cacioppo, J.T. (1986). The elaboration likelihood model of persuasion. *Advances in Experimental Social Psychology*, 19, 123-205.
- Python Reddit API Wrapper [Computer Software]. (2017). Retrieved from <https://github.com/praw-dev/praw>
- Stewart, B.D., & Payne, B.K. (2008). Bringing automatic stereotyping under control: Implementation intentions as efficient means of thought control. *Personality and Social Psychology Bulletin*, 34(10), 1332-1345.
- Strohming, N., & Nichols, S. (2014). The essential moral self. *Cognition*, 131(1), 159-171.
- Strohming, N., & Nichols, S. (2015). Neurodegeneration and identity. *Psychological Science*, 26(9), 1469-1479.
- Tan, C., Niculae, V., Danescu-Niculescu-Mizil, C., & Lee, L. (2016). Winning arguments: Interaction dynamics and persuasion strategies in good-faith online discussions. In *Proceedings of the 25th International Conference on World Wide Web* (pp. 613-624). Montreal, Canada.