Human Values and Trust in Scientific Journals, the Mainstream Media and Fake News

Nitin Verma
School of Information, University of Texas at Austin, USA.
nitin.verma@utexas.edu

Kolina S. Koltai
School of Information, University of Texas at Austin.
koltai@utexas.edu

Kenneth R. Fleischmann
School of Information, University of Texas at Austin, USA.
kfleisch@ischool.utexas.edu

ABSTRACT

What factors influence trust in online information? Americans increasingly get information from social media, public distrust in the mainstream media is growing and "fake news" is an important new phenomenon. This paper examines the factors that influence trust in scientific claims posted via social media, including the use of hyperlinks and readers' values. The paper describes a crowdsourcing-based experimental design using Amazon's Mechanical Turk platform. The core of the experiment was a set of 10 scientific findings reported in open-access, peer-reviewed scientific journals, which were in turn linked to in articles in both the mainstream media and "fake news" sites. Data analysis involved exploration of relationships between trust and the presence or absence of hyperlinks, and between trust and human values using nonparametric statistical methods. In terms of the influence of hyperlinks on trust, inclusion of hyperlinks to scientific journals, mainstream media articles, and even hidden URLs led to higher trust than hyperlinks to "fake news" sites or posts without hyperlinks ($p < 0.001$). Participants who clicked on hyperlinks to scientific articles placed higher trust in the claims than those who did not ($p < 0.001$). In terms of the influence of values on trust, values had the most impact in cases where individuals saw, but decided not to click on, hyperlinks; this finding seems to indicate that in the absence of firsthand examination of the hyperlinked sites, participants tend to rely more heavily on their values to determine their trust in a scientific claim. These findings indicate that both the presence and absence of hyperlinks and the values of the reader both significantly impact trust judgments.

KEYWORDS

human values, trust, online communication, social media, science communication, science journalism

80th Annual Meeting of the Association for Information Science & Technology, Washington, DC, VA | Oct. 27-Nov. 1, 2017

INTRODUCTION

In the 2016 U.S. presidential election, "fake news" stories outperformed mainstream media content on Facebook (Silverman, 2016). The prominent role of fake news on social media is alarming as about 14% of Americans identified social media as their "most important" source of election news (Allcott & Gentzkow, 2017). Berghel (2017) argues that a lack of emphasis on critical analysis in our educational system ill equips the general population to make informed judgments about what to trust. This raises the question whether such massive circulation effectively renders all information subjective (Hope, Hunter, & McLeod, 2015), ushering in a post-truth era (Harsin, 2015). At the same time, public trust toward institutions, particularly the mainstream media, has steadily declined (Smith & Son, 2013). According to Pew Research Center (Funk & Kennedy, 2016), while 76% of the public had either "a great deal" or "a fair amount" of confidence in scientists to act in the common good, only 38% shared that sentiment about the mainstream media.

Human values play an important role in influencing behavior, attitudes and sentiment (Fleischmann, 2014). Values can also predict attitudes toward opinion pieces about controversies (Templeton & Fleischmann, 2011). Price (1992) asserts that individuals rely on values, schemas, and group identifications to help guide their attitudes toward complex issues. This situation is particularly relevant in the social media era in which multiple sources of information compete for user attention and put pressure on users' capacity to vet all of the information that they consume.

The goal of our study was to examine the factors that influence trust in information, including how individuals with different life priorities, or human values, make different trust judgments toward information from scientific journals, the mainstream media and fake news. We set out to answer three main research questions. First, does the presence or absence of hyperlinks influence trust? Second, does the type of hyperlink used influence trust? Third, do the values of the reader influence trust? To answer these questions, we used an experimental design employing crowdsourcing, modeling how a Twitter user might review their timeline.

In this paper, we present results from our exploratory study of the relationships among hyperlinks, human values, and

trust in different media sources. In the following section, we discuss prior research on values and trust in online content. In the methods section, we describe the design of our study, including how we collected and analyzed the data. In the subsequent sections, we present our findings and discuss implications, limitations, and future research directions, followed by the conclusion of the paper.

BACKGROUND

Online social networks flatten hierarchies for publishing and information sharing, allowing anyone with an account to create, access, and distribute content. As a result, content gets published and shared without editorial oversight by subject-matter experts (Shao, Ciampaglia, Flammini, & Menczer, 2016). Traditionally, information originating from the mainstream media has been perceived to be more credible owing to the expectation that the organizations apply rigor and fact-checking before publishing (Flanagin & Metzger, 2007). On social media, fake news articles get shared alongside peer-reviewed scientific articles and fact-checked mainstream media articles. The many "economies of attention" thus generated within social networks compete for user attention and result in distinctive patterns of information production and user participation (Bessi, Scala, Rossi, Zhang, & Quattrociocchi, 2014). To understand the implications on trust in the different types of information sources, it is important to understand both the elements of trust in information (Kelton, Fleischmann, & Wallace, 2008) and users' motivations for using social media (Oh & Syn, 2015). Rieh (2014) compared users' evaluations of the trustworthiness of traditional media content and user generated content, finding that traditional media content was judged more trustworthy for health and news information.

Devos et al.'s (2002) work on trust towards institutions offers fertile grounds to argue that human values play a substantive role in formation of belief or knowledge about, and therefore trust towards, institutions. They found that values that emphasize stability, protection and preservation of tradition positively affect trust towards institutions, whereas values that emphasize independent thought affect such trust negatively (Devos et al., 2002). Whereas that study dealt with a broad set of institutions such as the healthcare system, the media, the law enforcement system, and political and religious institutions, our study focuses specifically on online media and how they report science news. It has also been argued that values, too, are influenced by media, and that people embed their values in both information and technology (Fleischmann, 2014). Existing research from a wide range of fields has produced multiple value inventories (reviewed in Cheng & Fleischmann, 2010); the most widely used value inventory is the Schwartz (1994) Value Inventory (SVI).

Gil & Artz (2007) offer a detailed set of criteria that people apply (consciously or subconsciously) to establish trust in online information. Their criteria include topic, context and criticality, popularity, authority, direct experience, recommendation, related resources, provenance, user expertise, bias, incentive, limited resources, agreement, specificity, likelihood, age, appearance, deception and recency. Subramaniam et al. (2015) add top-level domains, finding that users were generally more likely to trust .org sites than .com or .gov sites. Our study seeks to further build upon this research, examining in detail the influence of the presence of hyperlinks, the linked information itself, and the values of the reader on trust in online information.

METHODS

As mentioned in the introduction, we designed this exploratory study to look for relationships between values and trust in different types of information sources. At the outset, we wanted to create a study that closely approximated the information environment of our participants. This led us to two important design considerations: (i) the experiment had to be conducted with participants who had at least basic computer literacy and who were likely to be regular consumers of information online; and (ii) the core portion of the study instrument, where the participants interacted with articles, had to mimic a contemporary social media platform. Both of these design considerations led us to employ online crowdsourcing, as this approach allowed us access to an appropriate study population and the technical ability to set up a realistic Twitter timeline.

In the following subsections, we describe in detail our crowdsourcing platform, the process of building our corpus of scientific news claims, the design of our instrument, our data collection, our data cleaning, and our data analysis.

CROWDSOURCING PLATFORM

We used Amazon's Mechanical Turk (MTurk) as our crowdsourcing platform. We chose MTurk because comparative studies have found that MTurk performed favorably relative to alternatives such as university subject pools and posts on internet message boards in terms of accuracy of results and demographic diversity (e.g., Behrend, Sharek, Meade, & Wiebe, 2011; Hauser & Schwarz, 2016; Paolacci & Chandler, 2014). Specifically, MTurk has previously been applied to studies of trust in online information (e.g., Golbeck & Fleischmann, 2010) and the relationship between human values and attitudes toward news media texts (e.g., Templeton & Fleischmann, 2011). MTurk works by allowing "requesters" to design human intelligence tasks (HITs) and distributing them to "workers" on the platform. Workers then complete the HITs and receive monetary compensation from requesters.

SELECTION OF SCIENTIFIC NEWS CLAIMS

At the core of our HIT design was a corpus of scientific claims with corresponding scientific, mainstream media and fake news articles. The corpus was rigorously built through a process that involved multiple criteria. At the core of each post was a scientific claim that had to be published in a peer-reviewed academic journal and had to be openly accessible so that participants could read (or at least connect to and skim) the article for themselves. Thus, we selected only open-

access scientific articles. We limited the selection of peer-reviewed academic journals to those indexed in ISI Web of Science's Science Citation Index (SCI) Journal Citation Reports (JCR). We excluded articles that were corrected or retracted. As such, the set of open-access scientific articles were from high-quality venues and did not have any cloud of suspicion surrounding their publication.

In addition, each open-access scientific article had to be linked to by both a mainstream media outlet and a fake news site. We operationalized mainstream media using the list developed by Silverman (2016), referenced in Berghel (2017) and DiFranzo and Gloria-Garcia (2017); additionally, the selected sites were also independently identified by Flaxman, Goel, and Rao (2016) as among the 20 most popular news outlets. Specifically, the seven mainstream media outlets used in our corpus were *CBS News*, *CNN*, *Fox News*, *The Huffington Post*, *LA Times*, *USA Today*, and *The Washington Post*. The mainstream media article had to include a direct hyperlink to the open-access scientific article. The mainstream media articles also had to be native to the mainstream media outlets, rather than originating from another news service. The mainstream media articles had to be primarily textual in nature, in parallel with the open-access scientific articles, although some also had accompanying video clips. The mainstream media articles also had to appear on sites that were open-access without viewing restrictions. For example, we intentionally excluded *The New York Times* because it has a partial paywall that restricts the number of views per month for non-subscribers.

We operationalized fake news sites using a list from Fake News Watch (2016), previously employed by Shao, Ciampaglia, Flammini, and Menczer (2016). Specifically, two fake news sites used in our scientific claim corpus were InfoWars and Natural News. As in the case of the mainstream media articles, the fake news articles had to include direct hyperlinks to the open-access scientific articles, had to be native to the fake news site rather than originating from another source and had to be primarily textual in nature. This criterion was the primary reason why only two fake news sites were used; there is a great degree of cross-referencing among fake news sites, and several other fake news sites listed in the clickbait list from Fake News Watch are merely content aggregation sites, which do not produce their own original content.

The end result was a corpus of 10 scientific claims, listed in Table 1. Each claim originated from an open-access scientific article, which was directly hyperlinked-to by a mainstream media article and a fake news article. The carefully controlled development of this corpus makes it ideal for implementing in a HIT such as the one described in this paper.

HIT DESIGN

We designed our HIT using the TurkSuite Template Generator created by Alper (2014). The HIT comprised a series of 10 "screens" – the first three carried the informed consent form and instructions to complete the HIT, and the rest presented the various experimental and survey components of the study. In addition to collecting user-provided form data such as text input, Likert ratings, and drop-down selections, we recorded the amount of time a participant spent on each screen and whether or not they clicked on hyperlinks included in the HIT. We were sensitive to concerns of frequent exploitation of workers on MTurk (Irani, 2016), so we based our compensation on our estimated completion time – given our estimate of 20 minutes to complete the HIT, we offered $2.50 as compensation, equating to a pay rate above the Federal minimum wage. Fortunately, our results bore out that the average completion time was approximately 20 minutes, so, as such, we managed to comply with ethical compensation for our participants (Pittman & Sheehan, 2016). We also paid all workers who completed the HIT, whether or not we used their data in our final analysis, to reward their effort.

No.	Post Text
1	Gene editing tested for the first time in humans
2	Dolphins use specific whistles to call each other by a "name"
3	Children today now more prone to whooping cough because of weaker vaccine
4	Study found that GM Corn consumption by rats leads to organ damage
5	Some fast food packaging contains cancer-causing chemicals
6	Banning large sodas may lead to higher soda consumption
7	Human activity may lead to extinction of majority of primate species
8	Study links gut bacteria with chronic fatigue syndrome
9	Injecting mice brains with human fetal brain cells makes lab mice "smarter"
10	Ancient continental crust discovered in Indian Ocean under the island of Mauritius

***Table 1**. Scientific claims used in posts*

The core of our study was to ask participants to read a set of 10 posts as if they appeared in their social media news feed. Each of the10 posts was made up of two components. The first component was one of 10 scientific claims derived from the corpus described in the previous section. We attempted to keep the scientific claims easy to understand yet faithful to the original scientific article. To achieve this, the three authors each independently evaluated each set of articles (scientific, mainstream media and fake news) and independently proposed text for the scientific claim. The three authors then reviewed the independently generated text and reached consensus on the most appropriate text to use. We intentionally kept the scientific claims brief (8-12 words in length) to mimic social media content.

The second component of each post was a hyperlink. We created five experimental conditions using the hyperlink component as an independent variable. The first experimental condition involved not including any hyperlink. Each of the next three experimental conditions included a full hyperlink (i.e., the anchor text containing the full URL) to an open-access scientific article, a mainstream media article, or a fake news article. The fifth experimental condition used URLs shortened using Google's URL shortening service (Google, n.d.) to create hidden URLs (i.e., shortened URLs with the same generic base domain name). We refer to these different conditions as "post types" throughout the remainder of this article.

The aforementioned components informed the design of the emulated social media "newsfeed" component of the HIT. We decided to present our posts as tweets appearing on a typical Twitter user timeline. We mimicked the Twitter user timeline by emulating the HTML structure of individual tweets as rendered in a desktop web browser. In order to avoid potential confounds we did not emulate tweet features such as the number of replies, retweets, or likes since these could act as confounding cues for credibility. For similar reasons, we used generic handles (such as "Twitter User 1") and a neutral profile image (the Twitter new user "egg" avatar) to avoid any framing effects.

We displayed the emulated Twitter timeline on two screens: once by itself, and on the following screen with a 5-point Likert trust rating scale ranging from "Trust Completely" to "Don't Trust At All" against each individual post. We did this to separate the process of reading the posts, and of providing trust ratings. We intended this delay between the two phases to discourage workers from inattentively clicking on the radio buttons for ratings in order to complete the HIT quickly. Before displaying the first of these screens, we displayed a sample tweet screen to workers so as to familiarize them with what elements they would see on the emulated newsfeeds. We included the sample tweet screen to tell the users how to interact with the hyperlinks if they wished to. The posts appeared in identical order on both timeline screens. We recorded hyperlink clicking behavior (i.e., whether or not, and which, hyperlinks a worker clicked) independently for both timeline screens. Our need to record this clicking behavior was also one of the reasons for us to include the sample tweet screen described above. Due to limitations on tracking mouse clicks using Javascript, we advised workers to use only the left mouse click to open a hyperlink. At the same time, we provided them assurance that left-clicking the hyperlink would open the linked page in a new browser tab and would not interfere with the completion of their HIT.

One of our research objectives was to study whether human values influence people's trust in online news and their motivation to investigate hyperlinked resources in social media posts. To collect data on the values of our participants, we employed a gender-neutral version of the 21 item Portrait Values Questionnaire (PVQ) (Schwartz, 2003), adapted from the European Social Survey (ESS) (Schwartz, 2007). In the PVQ, each of the 21 statements provides a short personality "portrait," and participants indicate the extent to which the portrait describes their own personality. We used a 5-point Likert scale extending from "Absolutely Like Me" to "Not Like Me At All" for the participants to indicate the degree to which they believed each portrait described them. All 21 portraits were presented on the same screen, with the 5-point Likert scale appearing to the right of each item. Each participant was presented with the set of 21 portraits in a randomized order to mitigate any bias stemming from ordering or pattern effects. The 21 portraits correspond with Schwartz's ten value types (*power, achievement, hedonism, stimulation, self-direction, universalism, benevolence, tradition, conformity* and *security*), the complete middle-level of the Schwartz (1994) Value Inventory (SVI). We will discuss the process of calculating results from the level of 21 portraits to the level of ten value types in more detail below.

DATA COLLECTION

We created 100 HITs, each intended to be completed by a unique worker. We randomized the distribution of posts across the 100 HITs in such a way that each HIT had exactly two posts of a type, and that all five types were represented in each HIT for a total of 10 posts per HIT. We also ensured that all the 10 posts in a HIT were unique, implying that each HIT would carry 1 out of 5 versions per news story. We implemented the randomization and the subsequent apportioning of posts using a Python script.

As previously mentioned, we collected most of the data using the form elements in our HIT, but we also instrumented our HIT with Javascript to record the time workers spent on each screen, and which hyperlinks they clicked on.

DATA CLEANING

The 100 HITs were completed within approximately three hours of publication, which speaks both to the rapidity of data collection using MTurk and as an indication that our payment rate was deemed fair (at least relative to other HITs) by MTurk workers. We carefully reviewed all HITs for criteria such as uniqueness (we discarded one repeat HIT from a participant who completed the HIT twice), completeness (we

discarded eight HITs for being incomplete), correct answer to a dummy question (used as a validation step), and inconsistent answers to repeat questions (also used as a validation step; we discarded six HITs based upon this criterion), and inconsistent answers to questions about clicking on hyperlinks. For the last of the aforementioned criteria, we tallied worker responses to the following two assertions: "I carefully reviewed the linked articles" and "The content of the linked page influenced my trust in the post," with their actual clicking behavior. We found 16 people who had shown a general level of agreement with these assertions but had not clicked on *any* hyperlink. We defined "general level of agreement" as a "Strongly Agree" or the level just below and adjacent to it on the 5-point Likert scale. We reasoned that those who had not clicked on any hyperlink but reported that they had reviewed the content of the linked articles were contradicting themselves, and therefore their other responses could not be relied upon. The resulting dataset included complete and valid data from 69 unique participants.

DATA ANALYSIS

To facilitate analysis, the 5-point Likert scale corresponding to each portrait was converted to an ordinal scale ranging from -2 ("Not Like Me At All") to 2 ("Absolutely Like Me"). Then, for each of the 69 records, we converted the responses to the 21 PVQ portraits down to 10 values types using Schwartz's (2007) mapping of portraits to values types. Following Schwartz, while nine of the value types were measured using two portraits, one (universalism) was measured using three portraits. Schwartz justifies this additional acuity to measure universalism by citing that it is "the most complex value construct" (Schwartz, 2003). Following the calculation of the value type scores based upon Schwartz, we classified each participant as being "high" or "low" in terms of identification with a particular value. We defined "high" as a score higher than the median score for the value, and "low" as a score lower than or equal to the median score for that value.

Similarly, we converted the 5-point Likert scale for trust ratings into an ordinal scale ranging from -2 ("Don't Trust At All") to 2 ("Trust Completely"). Also, for this analysis, we combined the two independent sets of measurements of hyperlink clicking behavior into a third set: one binary value for each of the ten posts, indicating whether the hyperlink (if present) was clicked on or not. We achieved this by registering a particular hyperlink as clicked if it was clicked on in either of the two screens it was displayed on. If a hyperlink was not clicked on *either* screen, it was registered as not clicked. For each HIT, the number of hyperlinks clicked on was in the range 0 to 8 (since exactly two posts in each HIT were without any hyperlinks).

We used nonparametric methods to evaluate significance, as they carry fewer assumptions and are safer in cases where N is low (Gibbons & Chakraborti, 1991). We used Mann-Whitney U (alternatively, the Wilcoxon rank-sum test), the nonparametric equivalent to an unpaired t-test, and Kruskal-Wallis H, the nonparametric equivalent to ANOVA. Since the study was exploratory, all statistical tests were run two-tailed. We evaluated significance at the levels of $p < 0.05$, 0.01, and 0.001. To be conservative and reduce the chance of a Type I error, we focus most heavily on findings which were very highly significant, at the level of $p < 0.001$, as these significance levels were far less likely to occur randomly.

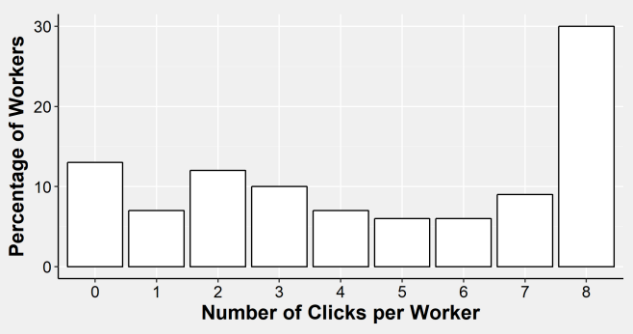

***Figure 1**. Distribution of the number of hyperlinks clicked by participants*

RESULTS

For this paper, we focused on the relationships among hyperlink clicking behavior, trust ratings, and human values across the five types of posts. We will begin with results related to hyperlink clicking behavior, and then move on to trust ratings, and then to human values.

HYPERLINK CLICKING BEHAVIOR

One preliminary finding across all participants ($N = 69$) was that the distribution of number of clicks per user ($M = 4.58$) had a relatively wide dispersion ($SD = 3.01$). As the bar graph in Figure 1 shows, 21 participants (30%) clicked on all eight hyperlinks in their HIT, while 9 participants (13%) did not click on any of the hyperlinks; the remaining 39 participants (57%) clicked on some but not all of the hyperlinks. Interestingly, Mann-Whitney U tests did not reveal any significant relationships between any of the ten values and the number of clicks per participant.

TRUST RATINGS BETWEEN POST TYPES

A Kruskal-Wallis H test for trust ratings across all five post types revealed a highly significant ($p < 0.001$) difference in trust ratings across post types. As the boxplots in Figure 2 show, trust ratings improved with the presence of hyperlinks. We then did Mann-Whitney U tests to compare each post type with every other post type. Based on the results, we can safely group the types into two groups, whereby trust in posts containing hyperlinks to scientific journals, mainstream media, and hidden URLs was higher than trust in fake news or posts without hyperlinks ($p < 0.001$ for all six pairwise comparisons). No significant differences were found in pairwise comparisons among the first three types, and hyperlinks to fake news earned more trust than posts without hyperlinks ($p < 0.05$).

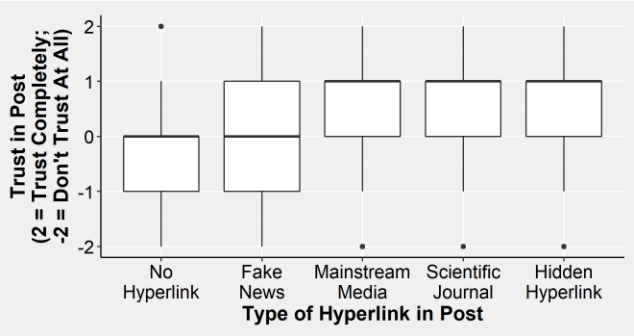

***Figure 2**. Boxplots of trust ratings across all five post types*

TRUST RATINGS AND HYPERLINK CLICKING BEHAVIOR (WITHIN POST TYPES)

Next, for the four post types that contained hyperlinks, we compared the trust ratings between individuals who did or did not click on the hyperlink to determine the effects of hyperlink clicking behavior. We found that clicking on a hyperlink significantly increased trust for posts with hyperlinks to fake news ($p < 0.05$), scientific journals ($p < 0.001$), and hidden URLs ($p < 0.001$). The boxplots for all four post types containing hyperlinks are shown in Figure 3. Plots on the left in each pair show the trust rating distribution when hyperlinks were not clicked, and those on the right show the distribution when hyperlinks were clicked.

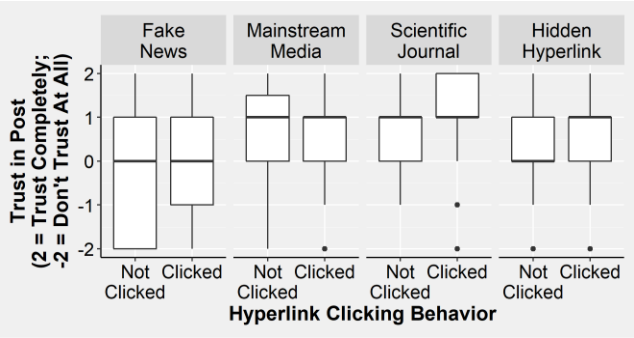

***Figure 3**. Trust ratings within each post type with clicking behavior*

TRUST RATINGS ACROSS POST TYPES (HYPERLINK NOT CLICKED)

For the case where the participant did not click on the hyperlink (Figure 4), Kruskal-Wallis H across all five post types) revealed a significant ($p < 0.001$) relationship between trust ratings and post type. To discern which post types were most significantly different, we ran ten Mann-Whitney U tests across pairs of post types. As shown in Figure 4, participants expressed significantly lower trust in posts without hyperlinks in comparison to posts with hyperlinks to mainstream media, scientific journals, or even hidden URLs ($p < 0.001$). Similarly, participants had significantly lower trust in posts with fake news hyperlinks than in posts with hyperlinks to mainstream media or scientific journals ($p < 0.001$).

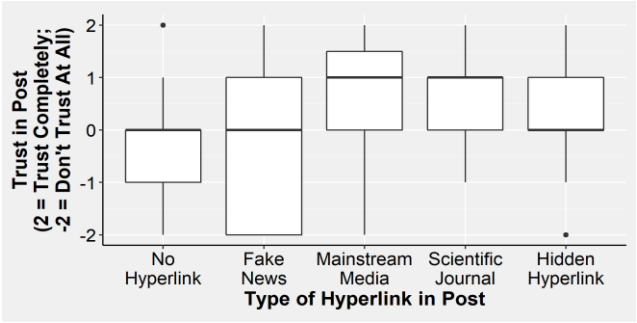

***Figure 4.** Trust ratings across post types (hyperlinks not clicked)*

TRUST RATINGS ACROSS POST TYPES (HYPERLINK CLICKED)

For the case where the participant clicked the hyperlink, Kruskal-Wallis H across all four post types containing hyperlinks revealed a highly significant ($p < 0.001$) relationship between trust ratings and post type. To discern which of the post types were significantly different, we ran pairwise Mann-Whitney U tests. As shown in Figure 5, participants had significantly lower trust in posts with fake news hyperlinks than in posts with hyperlinks to scientific journals ($p < 0.001$) or hidden URLs ($p < 0.001$).

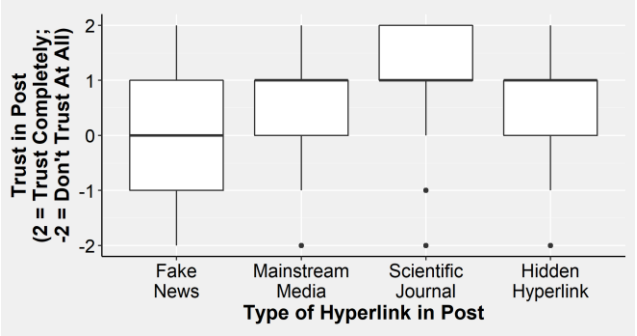

***Figure 5**. Trust ratings across post types (hyperlinks clicked)*

TRUST RATINGS AND VALUES

As explained previously, we reduced each worker's value scores for all ten SVI (Schwartz, 1994) values into a categorical variable with two levels: *high* (score above the median for that value) and *low* (score less than, or equal to, the median score). This categorization allows us to make assertions about what a person values (i.e., a "high" score), or does not value (i.e., a "low" score). We initially performed pairwise comparisons between the "high" and "low" for each of the ten values for each of the five post types. We found some statistically significant relationships, although none reached the level of $p < 0.001$. Given the differences that we had found between cases where individuals had clicked or had not clicked on a particular hyperlink, we performed more fine-grained comparisons, separating out the cases where the hyperlink was clicked or not clicked. We found more differences within the not clicked data than within the clicked data. Two of these differences were significant at the level of $p <$

0.001: highly valuing *stimulation* increased trust, while highly valuing *achievement* decreased trust. Considering the four conditions with hyperlinks which appear on both sides of the table, we identified a total of 15 significant differences; 13 of these differences were found in the not clicked condition (Table 2), while only two were found in the clicked condition (Table 3).

DISCUSSION

Overall, the results of our study demonstrate some interesting relationships among hyperlinks, trust, and human values. Specifically, we found that most participants clicked some, but not all, of the hyperlinks, illustrating that they made conscious choice about which hyperlinks to click. Participants demonstrated greater trust in posts that contained hyperlinks to scientific articles, mainstream media, and hidden URLs than to posts with no hyperlinks or hyperlinks to fake news, illustrating that participants were generally more willing to trust these more reputable sources despite the growing concerns about the influence of sharing of fake news via social media (Silverman, 2016). Trust was greater for scientific articles, hidden URLs, and fake news in cases where individuals clicked on hyperlinks; the first two results are unsurprising, particularly considering that they led to the same content; the latter result is interesting insofar as the content of the fake news appears to have outperformed the perceptions of the URL. However, due to the study design, it is not possible to tease out whether the greater trust resulted from reading the hyperlinked article or rather if it was greater trust that led to clicking the hyperlink in the first place, as these results indicate correlation rather than causation. Finally, values seemed to have a much more significant effect in cases where individuals did not click on the hyperlinks than when they did; this finding seems to add to the list of criteria for trust identified by Gil & Artz (2007), along the lines of Subramaniam et al. (2015) except going into attributes of the reader as well as the URL. This finding is particularly interesting and provides an opportunity for further research into how that trust in limited information such as URLs is mediated by the values of the reader.

LIMITATIONS

During the execution and analysis phases of our study, we identified limitations along the following themes:

Although we used a consensus based framing of the text of each post, the posts themselves dealt with only science-related topics. While this was a design choice to facilitate a focused exploration, the obvious limitation is that we cannot generalize our results for online posts outside the domain of science reporting.

Because we designed a "naturalistic" study (by selecting actual scientific news items) as against creating totally fictitious news items ourselves, there may be confounds (e.g., the lengths of the hyperlinked articles, writing styles, differing page layouts and ads) that could also affect trust.

Values ↓	No Hyperlink	Fake News	Mainstream Media	Scientific Journal	Hidden Hyperlink
Hedonism					
Stimulation				+	+ + +
Self-direction				+	
Universalism		- -	+ +	+	
Benevolence		- -			
Conformity		+			
Tradition	+ +			-	
Security					
Power			-		
Achievement			- - -	-	-

+: trust *increases* as importance of value *increases*
-: trust *decreases* as importance of value *increases*
+++/- - -: $p < 0.001$; ++/--: $p < 0.01$; +/-: $p < 0.05$

Table 2*. Relationship between trust ratings and importance of values (hyperlinks not clicked)*

Values ↓	Fake News	Mainstream Media	Scientific Journal	Hidden Hyperlink
Hedonism				
Stimulation				
Self-direction				
Universalism				
Benevolence				
Conformity		-		
Tradition	+			
Security				
Power				
Achievement				

+: trust *increases* as importance of value *increases*
-: trust *decreases* as importance of value *increases*
+++/- - -: $p < 0.001$; ++/--: $p < 0.01$; +/-: $p < 0.05$

Table 3*. Relationship between trust ratings and importance of values (hyperlinks clicked)*

Our sample size is relatively small ($N = 69$), and thus it is possible that we obtained false negative findings because we did not have sufficient power based on the number of participants used. Given that we used very conservative non-parametric analyses, it is far less likely that we had false positive results, so as such, a larger N would possibly reveal additional significant findings, but is less likely to invalidate any of the significant findings from this study, particularly those at the level of $p < 0.001$.

Ours was a "laboratory" experiment, conducted in a simulated environment where there is an existing "marketplace" of workers with predetermined skill sets, rather than an observational field study of users' exploration of and trust in news content within their own social media accounts.

Our experimental design does not allow us to infer causality in the decision to click or not click on a hyperlink; for example, when comparing individuals who clicked or did not click on a given hyperlink, the differences could be explained by participants learning from clicking and changing their degree of trust as a result, or by participants with more (or less) initial trust based on the hyperlink feeling more (or less) inclined to click on a hyperlink. Potentially, we could ask users to look at a URL first, measure trust, ask them to visit the website, and measure trust again; this type of study would be best performed in a usability lab setting.

FUTURE RESEARCH

One future direction would be to replicate the current study with a larger *N*, which is our next planned step. The study design could also be applied to domains beyond science, such as politics or health information. During the process of developing the corpus of posts we observed that fake news sites appeared to include hyperlinks to open-access scientific articles more often than their mainstream media outlet counterparts; systematic examination of this phenomenon was beyond the scope of this project but would make interesting further research. We also believe that seeking more incisive inputs from participants about the heuristics they deploy in establishing trust online would provide further nuance to our findings, such as being able to determine causality in the relationships we have discovered. Future efforts should also seek to minimize dependence on participant-reported data about whether, and to what extent, they examined the hyperlinked articles. One way to do this would be to conduct follow-up studies in a controlled environment such as a usability lab. Motivated by our findings pertaining to the influence exerted by values on people's trust judgment, the role played by the type of media outlet in persuading people, and the body of literature on trust, we encourage others to conduct similar studies, and therefore add to the research on trust in online information.

CONCLUSION

The dramatic rise in the volume of content posted online threatens to render most information "subjectively true" (Hope et al., 2015), and seems to be creating "truth markets" (Harsin, 2015) that make the task of establishing trust very demanding. While the Web, in facilitating the creation of this content, has created an egalitarian society when it comes to information access and diffusion, much of this content gets published without any editorial or journalistic oversight. In addition, the sheer volume and speed of this information puts pressure on people's limited attentional resources when they are confronted with the task of making a trust judgment. Given this complex situation, there is a need to understand the factors that influence trust in information exchanged via social media.

To develop that understanding, our study explored the role of hyperlinks and human values in trust in information exchanged via social media. Some of our findings were in line with our expectations. For example, we had expected that the presence or absence of a hyperlink in a post would influence trust in that post. We found that in comparison to the condition where there is no hyperlink, each post in which there was some type of hyperlink (fake news, mainstream, scientific journal, or even a hyperlink to a hidden URL) had a significantly *higher* trust level.

In addition, we found that the different types of hyperlinks elicited significantly different levels of trust, with trust in posts with hyperlinks to scientific journal, mainstream media, and hidden URLs higher than in posts with fake news hyperlinks or no hyperlinks. We discovered that clicking or not clicking (and thus, presumably, inspecting or not inspecting the contents of) a hyperlink significantly affected trust ratings only for posts carrying hyperlinks to scientific journals and hidden URLs, and to a marginally significant extent for posts with fake news hyperlinks. Specifically, posts with these three types of hyperlinks fared better in trust ratings when the corresponding hyperlinks were clicked on. We cannot assume with absolute certainty that mere clicking or opening of a hyperlink implies assessment of the contents of the linked page. Viewing scientific journal or fake news articles was correlated with higher trust ratings relative to choosing not to view them. Perhaps, for the case of scientific journals, this is because participants were not previously familiar with the outlets, but were impressed by the format and formality of the journal web pages. Further, for fake news articles, participants may have been influenced positively by seeing hyperlinks to the scientific journals (see footnote 5), or it is also possible that individuals who immediately had low trust in the fake news sites based upon the hyperlinked URLs chose not to click on them; determining the reason would be an interesting opportunity for future research. Fourth, among posts in which the embedded hyperlinks were clicked, we found that fake news was less trusted than the other outlets.

For values, the two most significant relationships involved the values of *achievement* and *stimulation*. First, when participants were shown but chose not to click the hyperlink to a mainstream media article, individuals who highly valued *achievement* were less likely to trust mainstream media articles than those who put less value on *achievement*. Schwartz (2007) defines *achievement* as "personal success through demonstrating competence according to social standards" (p. 174), so it seems reasonable to conclude that individuals who highly value *achievement* are more likely to care about social standards of excellence (and therefore seek to distinguish themselves from others), and are thus less inclined to trust mainstream media sources that are adopted by the masses. Second, when the participants were shown but chose not to click on the hyperlink to a hidden URL, individuals who highly valued *stimulation* were more likely to trust the article pointed to by the hidden URL than those who put less value on *stimulation*. Schwartz (2007) defines *stimulation* as "excitement, novelty, and challenge in life" (p. 174), so it seems reasonable to conclude that individuals who highly value *stimulation* are more likely to take risks, such as blindly trusting a hyperlink with a hidden URL. Overall, the results for

participants who did not click the hyperlink were more influenced by values than those who did click the hyperlink; one possible explanation is that when shown a hyperlink, in the absence of more information, individuals make gut trust decisions based upon their values; however, when individuals explore the linked content, they have other factors that they can use to inform their trust judgments.

Thus, this study provides preliminary insights about the role that human values play in influencing trust in different media outlets. We believe that further research is called for to better understand how human values shape trust in online content.

ACKNOWLEDGEMENTS

The authors wish to acknowledge funding from a Special Research Grant from the Office of the Vice President for Research of the University of Texas at Austin, as well as the anonymous crowdworkers who participated in this study.

REFERENCES

Allcott, H., & Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, *31*, 211-236.

Alper, M. (2014). TurkSuite Template Generator. Retrieved April 23, 2017, from http://mturk.mit.edu/template.php

Behrend, T. S., Sharek, D. J., Meade, A. W., & Wiebe, E. N. (2011). The viability of crowdsourcing for survey research. *Behavior Research Methods*, *43*, 800-813.

Berghel, H. (2017). Lies, damn lies, and fake news. *Computer*, *50*(2), 80-85.

Bessi, A., Scala, A., Rossi, L., Zhang, Q., & Quattrociocchi, W. (2014). The economy of attention in the age of (mis)information. *Journal of Trust Management*, *1*, 12.

Cheng, A.-S., & Fleischmann, K. R. (2010). Developing a meta-inventory of human values. *Proceedings of the American Society for Information Science and Technology*, *47*, 1-10.

Devos, T., Spini, D., & Schwartz, S. H. (2002). Conflicts among human values and trust in institutions. *British Journal of Social Psychology*, *41*(4), 481-494.

DiFranzo, D., & Gloria-Garcia, K. (2017). Filter bubbles and fake news. *XRDS*, *23*(3), 32-35.

Fake News Watch. (2016, January 18). Retrieved April 23, 2017, from http://fakenewswatch.com/

Flanagin, A. J., & Metzger, M. J. (2007). The role of site features, user attributes, and information verification behaviors on the perceived credibility of web-based information. *New Media & Society*, *9*, 319-342.

Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, *80*, 298-320.

Fleischmann, K. R. (2014). *Information and human values*. San Rafael, CA: Morgan & Claypool.

Funk, C., & Kennedy, B. (2016). *The politics of climate*. Pew Research Center.

Gibbons, J. D., & Chakraborti, S. (1991). Comparison of the Mann-Whitney, Student's t, and alternate t tests for means of normal distributions. *The Journal of Experimental Education*, *59*, 258-267.

Gil, Y., & Artz, D. (2007). Towards content trust of web resources. *Web Semantics: Science, Services and Agents on the World Wide Web*, *5*(4), 227-239.

Golbeck, J., & Fleischmann, K. R. (2010). Trust in social Q&A: The impact of text and photo cues of expertise. *Proceedings of the American Society for Information Science and Technology*, *47*, 1-10.

Google. (n.d.). Google URL shortener. Retrieved April 23, 2017, from https://goo.gl/

Harsin, J. (2015). Regimes of posttruth, postpolitics, and attention economies. *Communication, Culture & Critique*, *8*(2), 327-333.

Hauser, D. J., & Schwarz, N. (2016). Attentive Turkers: MTurk participants perform better on online attention checks than do subject pool participants. *Behavior Research Methods*, *48*(1), 400-407.

Hope, A., Hunter, R., & McLeod, J. (2015). Climate change and social media: Trust, literacy, legitimacy and subjectivity. Presented at the Global Sustainability Institute Research Conference 2015, Cambridge, UK. Retrieved from http://nrl.northumbria.ac.uk/21360/

Irani, L. (2016). The hidden faces of automation. *XRDS*, *23*(2), 34-37.

Kelton, K., Fleischmann, K. R., & Wallace, W. A. (2008). Trust in digital information. *Journal of the American Society for Information Science and Technology*, *59*, 363-374.

Oh, S., & Syn, S. Y. (2015). Motivations for sharing information and social support in social media: A comparative analysis of Facebook, Twitter, Delicious, YouTube, and Flickr. *Journal of the Association for Information Science and Technology*, *66*, 2045-2060.

Paolacci, G., & Chandler, J. (2014). Inside the Turk: Understanding Mechanical Turk as a participant pool. *Current Directions in Psychological Science*, *23*, 184-188.

Pittman, M., & Sheehan, K. (2016). Amazon's Mechanical Turk a digital sweatshop? Transparency and accountability in crowdsourced online research. *Journal of Media Ethics*, *31*(4), 260-262.

Price, V. (1992). *Public opinion*. Newbury Park, CA: Sage.

Rieh, S. Y. (2014). Credibility assessment of online information in context. *Journal of Information Science Theory and Practice*, *2*(3), 6-17.

Schwartz, S. H. (1994). Are there universal aspects in the structure and contents of human values? *Journal of Social Issues*, *50*(4), 19-45.

Schwartz, S. H. (2003). A proposal for measuring value orientations across nations. *Questionnaire Package of the European Social Survey*, 259-290.

Schwartz, S. H. (2007). Value orientations: Measurement, antecedents and consequences across nations. In R. Jowell, C. Roberts, R. Fitzgerald, & G. Eva (Eds.), *Measuring attitudes cross-nationally: Lessons from the European Social Survey* (pp. 169-203). London: Sage.

Shao, C., Ciampaglia, G. L., Flammini, A., & Menczer, F. (2016). Hoaxy: A platform for tracking online misinformation. In *Proceedings of the 25th international conference companion on world wide web* (pp. 745-750). Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee. https://doi.org/10.1145/2872518.2890098

Silverman, C. (2016, November 16). This analysis shows how viral fake election news stories outperformed real news on Facebook. Retrieved April 21, 2017, from https://www.buzzfeed.com/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook

Smith, T. W., & Son, J. (2013). General Social Survey 2012 final report: Trends in public attitudes about confidence in institutions. NORC at the University of Chicago.

Subramaniam, M., St. Jean, B., Taylor, N. G., Kodama, C., Follman, R., & Casciotti, D. (2015). Bit by bit: Using design-based research to improve the health literacy of adolescents. *JMIR Research Protocols*, *4*(2), e62.

Templeton, T. C., & Fleischmann, K. R. (2011). The relationship between human values and attitudes toward the Park51 and nuclear power controversies. *Proceedings of the American Society for Information Science and Technology*, *48*, 1-10.