



Community Fact-Checks Trigger Moral Outrage in Replies to Misleading Posts on Social Media

Yuwei Chuai
University of Luxembourg
Luxembourg, Luxembourg
yuwei.chuai@uni.lu

Gabriele Lenzini
University of Luxembourg
Luxembourg, Luxembourg
gabriele.lenzini@uni.lu

Anastasia Sergeeva
University of Luxembourg
Luxembourg, Luxembourg
anastasia.sergeeva@uni.lu

Nicolas Pröllochs
JLU Giessen
Giessen, Germany
nicolas.proellocks@wi.jlug.de

Abstract

Displaying community fact-checks is a promising approach to reduce engagement with misinformation on social media. However, how users respond to misleading content *emotionally* after community fact-checks are displayed on posts is unclear. Here, we employ quasi-experimental methods to *causally* analyze changes in sentiments and (moral) emotions in replies to misleading posts following the display of community fact-checks. Our evaluation is based on a large-scale panel dataset comprising $N = 2,225,260$ replies across 1841 source posts from X's Community Notes platform. We find that informing users about falsehoods through community fact-checks significantly increases negativity (by 7.3%), anger (by 13.2%), disgust (by 4.7%), and moral outrage (by 16.0%) in the corresponding replies. These results indicate that users perceive spreading misinformation as a violation of social norms and that those who spread misinformation should expect negative reactions once their content is debunked. We derive important implications for the design of community-based fact-checking systems.

CCS Concepts

• **Human-centered computing** → **Social media; Empirical studies in collaborative and social computing**; • **Information systems** → **Crowdsourcing**; • **Security and privacy** → **Human and societal aspects of security and privacy**.

Keywords

Misinformation, fact-checking, social media, crowdsourcing, online emotions, moral outrage

ACM Reference Format:

Yuwei Chuai, Anastasia Sergeeva, Gabriele Lenzini, and Nicolas Pröllochs. 2025. Community Fact-Checks Trigger Moral Outrage in Replies to Misleading Posts on Social Media. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 23 pages. <https://doi.org/10.1145/3706598.3713909>



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '25, Yokohama, Japan*

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1394-1/25/04
<https://doi.org/10.1145/3706598.3713909>

1 Introduction

Social media platforms like X (formerly Twitter), Facebook, and TikTok have become increasingly influential sources for news consumption and social engagement, especially among younger generations [31, 117]. However, the fast-paced and interactive nature of social media has also compromised the quality of the information shared, as content can be posted rapidly without undergoing the thorough fact-checking essential to journalistic ethics [85]. As a result, misinformation spreads easily on social media platforms and has become a concerning issue in the digital age [48, 70, 121]. The negative effects of user engagement with online misinformation have been repeatedly observed across various domains [1, 37, 38, 42, 52, 90]. For instance, political misinformation is especially troubling, as engagement with false information during elections can undermine democracy and destabilize societies [6, 37, 52, 121]. Given the significant risks and threats posed by online misinformation, social media providers are thus urged to establish content moderation policies and implement effective interventions to identify misinformation and inform users on their platforms [15, 32, 34, 70].

As a countermeasure, the social media platform X/Twitter has recently adopted a crowdsourced (community-based) fact-checking approach, “Community Notes,” to fact-check misleading posts and communicate veracity to users on its platform [94]. “Notes” represents the first attempt by a major social media platform to apply a crowdsourced approach to fact-checking on a large scale. It allows users to add context or correct information to posts via *community notes*, providing additional perspectives or clarifications. Once community notes are rated helpful by other community contributors, they will be directly displayed on the corresponding misleading posts. While the accuracy of crowdsourced fact-checking approaches has been thoroughly validated by previous research [8, 33, 43, 89, 103], a causal understanding of how users react to social media content that has been subject to community-based fact-checking is still in its infancy.

Previous research studying the effect of community fact-checks on social media users has primarily focused on analyzing how it affects *engagement* with misinformation (e. g., shares, comments, replies, or likes) [20, 21, 33, 36, 97]. For instance, studies have shown that community notes can reduce users’ belief in false content and their intentions to share misleading posts [19, 20, 35, 97]. However, little is known about how social media users *emotionally* react to social media posts that have been subject to fact-checking. Emotions

are highly influential for human decision-making and judgment [72]. It is thus conceivable that fact-checking social media posts triggers emotional backlash towards the authors of such content.¹ For example, users may feel a sense of urgency to correct the record or call out the spreaders of misinformation. In this scenario, the debunking of misinformation via community notes might provoke strong emotional reactions (e. g., negative sentiment, anger, disgust) from social media users. Note that such emotional backlash from the user base may be seen as a double-edged sword. On the one hand, it can be a desirable reaction in the fight against misinformation as it may signal to the content creator and other users that spreading misinformation is unacceptable and subject to public scrutiny. Also, collective emotional backlash may help pressure authors to correct or delete misleading content and discourage others from sharing or engaging with it. On the other hand, however, social media platforms also need to maintain a healthy and constructive dialogue. In this regard, the emotional reactions from the user base pose risks of fueling political polarization, escalating conflicts, and entrenching divisions. Yet, empirical evidence on how fact-checking emotionally resonates with the audience on social media is missing.

Furthermore, fact-checking social media posts may trigger pronounced moral reactions from the user base. Content moderation of online speech is a moral minefield, especially when two key values come into conflict: upholding freedom of expression and preventing harm caused by the violation of social norms and spreading misinformation [68]. From a theoretical perspective, spreading online misinformation may be perceived as immoral behavior, motivating users' moral reasoning and judgment and triggering their moral emotions [84]. Previous work has identified the phenomenon of *moral outrage*, which occurs when individuals perceive the actions of another person or an entire situation as a serious violation of their moral principles or ethical standards. Moral outrage is commonly associated with a higher likelihood of inter-group hostility and can contribute to the further escalation of conflicts on the platforms [28, 84, 117]. Emotionally, this often results in a combination of anger and disgust, prompting individuals to express these emotions in an attempt to restore what they perceive as the proper moral norm [104]. This response typically involves blaming the violator and demanding corrective actions [84, 117]. Based on this rationale, we hypothesize that social media users might react with expressions of moral outrage to misleading posts and their authors once they are informed about the falsehoods via community notes.

Research questions: In this study, we *causally* analyze whether the display of community notes changes expressions of sentiments and (moral) emotions in social media users' replies to misleading posts. Specifically, we address two main research questions. Our first research question is to examine the effects of community notes on sentiments and basic emotions in replies (RQ 1.1). Previous research further suggests that emotions may play a particularly pronounced role in the context of political (mis-)information [11,

95]. We thus additionally analyze how the emotional reactions differ across misleading posts that cover a political vs. a non-political topic (RQ1.2):

- **RQ1.1:** Which sentiments (positivity, negativity) and basic emotions (e. g., anger, disgust) in replies to misleading posts are triggered by the display of community notes?
- **RQ1.2:** Do the effects of community notes on sentiments and basic emotions in replies differ across political and non-political misleading posts?

Furthermore, when anger and disgust are mixed, they generate “moral outrage,” a more intense emotional reaction to perceived moral transgressions than either anger or disgust alone [9, 28, 104]. Our second research question is to examine whether community notes trigger moral outrage in replies (RQ2.1), and analyze the sensitivity across political and non-political misleading posts (RQ2.2):

- **RQ2.1:** Does the display of community notes trigger moral outrage in replies to misleading posts?
- **RQ2.2:** Do the effects of community notes on moral outrage in replies differ across political and non-political misleading posts?

Data & methods: To address our research questions, we collect a dataset of $N = 2,225,260$ replies to 1841 misleading posts that have been fact-checked via community notes over a 4-month period since the roll-out of “Community Notes” feature on X/Twitter. The dataset includes all the replies that are directly toward misleading posts before and after the display of community notes. Subsequently, we employ methods from Natural Language Processing (NLP) in combination with Regression Discontinuity Design (RDD) to estimate the causal effects of community notes on the sentiments and (moral) emotions in the replies to the fact-checked posts.

Contributions: To the best of our knowledge, this study is the first to causally analyze the emotional reactions of social media users in response to (community) fact-checks. We find that informing users about falsehoods through the display of community fact-checks significantly increases the expression of negativity in the corresponding replies by 7.3%. Furthermore, we provide a fine-granular analysis of emotions. Here, the largest effect sizes are observed for (negative) emotions from the moral emotions family. Specifically, displaying community notes increases anger in the corresponding replies by 13.2%, disgust by 4.7%, and moral outrage by 16.0%. The observed effects are significantly more pronounced for political vs. non-political posts. Overall, these results imply that social media users perceive misinformation as moral transgressions and that authors of misinformation should expect negative reactions once their content is debunked. Based on our findings, we derive important implications for behavioral theory, media literacy, and the design of community-based fact-checking systems on social media.

2 Background and Related Work

2.1 Misinformation on Social Media

Social media platforms, as a source for news consumption and social engagement, have become increasingly popular. As of July 2024, there is a large population base of ~5.17 billion users on social media, and the number continues to grow [113, 117]. Especially among younger users aged 18 to 24, up to 65% prefer social media

¹Note that an alternative rationale posits that some users (e. g., those with strong partisan alignments) may respond emotionally to a post from their “side” being labeled as misleading, rather than reacting to the misinformation itself. However, previous research suggests that replies are typically aimed at the content or the post authors [129]; and users are more likely to reply to authors from ideologically opposing groups [49]. In line with this notion, our later analysis and manual validation imply that community notes primarily elicit replies focused on the content of misleading posts or their authors (see details in Section 4.4).

as one of their preferred news sources [31]. Together, these data indicate a trend toward partially substituting traditional media with social media channels that provide fast access to information [12] and offer a more interactive way to engage with content [105].

Due to the lack of editorial oversight, social media platforms have also become fertile soil for spreading misinformation. Recently, there have been some critics arguing that the prevalence of misinformation on social media is overestimated, and the exposure is concentrated among narrow fringe groups [4, 13]. However, given the large population base on social media platforms, its negative effects can be profound. For example, the spread of online health misinformation has posed severe risks to public health in both the Global North and the Global South [42, 65, 90, 120]. Additionally, the rise of AI makes generating misinformation faster, which poses new challenges and further undermines the credibility of social media [45, 50, 130, 132].

The spread of political misinformation on social media:

The spread of political misinformation on social media has become a tool to pursue political agenda, and its negative effects are continually observed in threatening the stability of societies and eroding trust in democracy [1, 37, 38]. For example, conspiracy theories claiming that the 2020 U.S. presidential election was “stolen” were widely shared on social media. It is found that users who were exposed to at least one misinformation website are 17.3% more likely to believe that the election was fraudulent [29]. As surveyed, almost 40% of the U.S. public and, especially, almost 70% of Republicans negated the legitimacy of the 2020 presidential election outcome [112]. The reduced faith in electoral institutions have had real-world consequences in affecting voters’ behaviors in subsequent elections and hate campaigns [37, 52]. Similar patterns have been observed in countries of the Global South, such as in Brazil [54, 96]. Moreover, political misinformation is often interconnected with the politicization of other social issues, such as health crises and climate change, further contributing to polarized attitudes toward these topics [23, 24, 55].

Altogether, with the increasing use of social media for news consumption and the continuously observed negative effects of misinformation across various domains, the spread of misinformation on social media has become a significant societal issue in the digital age [117, 121]. Particularly, political misinformation has been found to spread significantly faster than non-political misinformation on social media [121]. Previous studies also show that social media users share information to single their political affiliation and moral stance, even knowing that it is inaccurate [77, 86]. To protect users from being affected by misinformation and prevent harm, it is both warranted and required to implement effective countermeasures to fight against its spread [6, 37].

2.2 Content Moderation for Misinformation

Social media providers are increasingly urged to establish content moderation policies to regulate online behavior and combat the spread of misinformation on their platforms [32, 70]. The content moderation process includes two main aspects: fact-checking misinformation and handling identified misinformation.

To identify misinformation, social media platforms are widely collaborating with professional third-party fact-checkers (e. g., the

International Fact-Checking Network) [78]. While fact-checking experts can accurately access veracity [71, 79], the implementation of this approach faces two main challenges. First, professional fact-checking is a lengthy and resource-consuming process, leading to a situation where only a small proportion of posts are actually checked [76, 79, 91, 123]. Second, professional fact-checkers face suspicions of (political) bias, e. g., by selectively fact-checking content [24, 116]. Therefore, the trust of the public in professional fact-checking is discounted [19, 36, 47]. To overcome these challenges, researchers have proposed to fact-check and identify misleading posts on social media platforms at scale based on the non-expert fact-checkers from the crowd [2, 80, 89]. While the individual assessments can have bias and noise [3], it has been shown that the accuracy of aggregated judgments, even from relatively small crowds, is reliable and comparable to the accuracy of expert fact-checkers [8, 43, 89, 103]. Additionally, the crowdsourced assessments are perceived as more trustworthy by the users than fact-checks from experts [2, 36, 131].

Subsequently, social media platforms generally have three options for handling identified misleading posts: removing them, reducing their visibility, or informing users about the falsehoods [64, 78, 124]. The first two options – removal and visibility reduction – are forms of “hard moderation,” where the platforms limit or entirely cut off the spread of misinformation. In contrast, informing other users that a post is misleading represents “soft moderation,” where the platforms refrain from imposing direct controls and instead allow users to make their own decisions on how they process the informed misleading post. The hard moderation approaches face a dilemma, where the two key values of protecting freedom of expression and preventing harm from the spread of misinformation come into conflict [68, 107]. Given this, soft moderation approaches can be regarded as a trade-off between freedom of expression and misinformation intervention. Previous lab/survey experiments consistently show that communicating the falsehood to users (especially with explanations) is effective in reducing users’ beliefs and sharing intentions [2, 26, 57, 59, 62, 92, 127, 131].

2.3 Sentiments, Basic Emotions, and Moral Outrage on Social Media

As an essential element of human communication, sentiments and emotions play a pivotal role in shaping interactions on social media platforms [7, 63, 72, 114]. They can be understood as two affective dimensions: valence (sentiment) and arousal (emotion). Sentiment, i. e., the valence of content (positive vs. negative), is an attitude, thought, or judgment promoted by a feeling [126]. Research has shown that negativity is a particularly significant characteristic of online engagement [30, 99, 133]. Negative sentiment tends to evoke stronger reactions, promoting users to comment and share online posts, especially among polarized communities [30, 133].

The relationship between sentiments and social interaction is more nuanced than simply the valence (positive vs. negative) of the content alone. From the arousal dimension, specific emotions (i. e., affect ladenness) are embedded within the content, adding intensity to the positive or negative feelings expressed [87, 126]. With various models classifying basic emotions, Ekman’s six basic emotion model is widely used and commonly shared across

other frameworks [126]. Specifically, Ekman’s model identifies six universal emotions: anger, disgust, fear, joy, sadness, and surprise [40, 41, 106]. These emotions are considered fundamental across different cultures and can combine to form more complex emotional experiences [106]. Emotions also vary on the level of physiological arousal or activation they evoke [7]. For example, anger and sadness are often expressed as negative emotions. However, anger is a state of heightened arousal and makes online content more viral, while sadness is low arousal and makes online content less viral [7, 18]. Particularly, users’ reliance on emotions can promote their belief in misinformation [75]. Several studies have found that emotions can explain the spread of true vs. false information on social media, with anger being the most significant emotion in driving the spread of misinformation [22, 93].

Additionally, *moral emotions* can arise in response to perceived moral violations or standards and motivate moral behavior [53]. One large family of moral emotions is other-condemning emotions that often lead to behavioral outcomes, that is, actions to correct the other’s behavior and increase justice-relevant judgment [53, 118]. Specifically, anger and disgust are the main other-condemning emotions used in coordinating and constraining the behavior of the group member who made the morally unacceptable actions [53]. During the COVID-19 pandemic, other-condemning emotions are significantly associated with the higher virality of misinformation on social media [111]. Within political contexts, social media posts containing moral-emotional language have been shown to be 20% more likely to be shared, particularly within polarized partisan networks (liberal or conservative), with stronger effects within these groups than between them [11]. This highlights the role of moral emotions in amplifying online engagement, particularly through expressions of out-group animosity. Such expressions, especially within partisan networks, have been shown to drive higher engagement than in-group favoritism, further entrenching political polarization [95].

Previous studies highlight *moral outrage* – a mixture of anger and disgust triggered by a perceived moral norm violation – as a powerful emotion that motivates people to shame and punish wrongdoers, especially in the digital age [10, 28, 104, 110]. Social media platforms, such as X/Twitter, inherently amplify expressions of moral outrage over time, as users who learn such language get rewarded with an increased number of likes and reposts [9]. Moral outrage is also contagious. Knowing that a person or entity is being publicly condemned can provoke moral outrage in others [108]. This contagious has the potential to ignite online firestorms – large waves of insults, criticism, or swearing directed at persons and organizations on social media [101, 115]. In addition, users frequently overestimate the level of outrage and hostility in social media posts, perceiving emotional expressions as more extreme than the authors actually report. This overperception can lead to escalated responses and increase the spread of moral outrage [10]. The consequences of heightened moral outrage on social media are double-edged. On the one hand, it can fuel inter-group conflict and polarization, exacerbating online tensions and divisions. On the other hand, it can also amplify constructive facets of morality, such as social support and prosociality [28, 117]. This dual role underscores the importance of understanding how moral outrage

operates in specific circumstances to mitigate its negative impacts while promoting its potential positive effects.

Overall, the analysis of sentiments and (moral) emotions is a crucial and growing area of research for understanding online engagement and its consequences [46]. A large research body has documented that sentiments, (moral) emotions, and moral outrage play a pronounced role in shaping engagement with (mis-)information on social media. However, an understanding of how users emotionally react to social media posts after they are informed of the falsehoods is missing, which presents the aim of our study.

3 Data and Methods

3.1 Datasets

In this study, we analyze the emotional reactions of social media users in response to (community) fact-checks. For this purpose, we collect a large dataset of posts that have been fact-checked on X’s (formerly Twitter) “Community Notes” platform (formerly Birdwatch). “Community Notes” was officially rolled out to users on December 11, 2022, and is a crowdsourced fact-checking program that allows users to identify posts that they believe are misleading and write short textual notes that provide context to the posts. After a community note is created, other users can rate its helpfulness and, if it reaches a certain level of helpfulness, it is displayed prominently beneath the original post. An example of a misleading post with a displayed community note and one direct reply is shown in Fig. 1a. Importantly, before note display, users who reply to the source posts can only see the original post content (as shown in Fig. 1b). After note display, users who reply to the source posts can see both original post content and displayed community notes flagging the post content as being misleading.

As detailed in the following, we collect data from three different sources: (i) note dataset, (ii) post dataset, and (iii) reply dataset. Subsequently, we combine these three data sources into a longitudinal dataset that allows us to analyze the changes in reply emotions before and after the display of community notes.

3.1.1 Note dataset. X publishes all community notes and their associated status histories on a dedicated website and updates them daily.² We downloaded the community note datasets on April 8, 2023, which include all note contributions up until April 6, 2023. Our note dataset contains the details for each note, including its unique ID, creation time, classification (misleading or not misleading), referenced source post ID, and the note status history (i. e., when the notes were rated helpful and displayed on the source posts). In our data, there are 68,235 community notes for 45,664 source posts. However, only 4686 source posts (10.3%) received displayed community notes until the date of data collection.

3.1.2 Post dataset. We use X’s Academic Research API to search for the 4686 source posts with displayed community notes and successfully collect 3872 (82.6%) posts. Notably, given that the “Community Notes” program was rolled out on December 11, 2022, and mainly developed in the U.S., we only consider the source posts that were created after the roll-out and are written in English. There are 1953 source posts remaining after this filtering step.

²All community notes are publicly available at <https://x.com/i/communitynotes/download-data>.

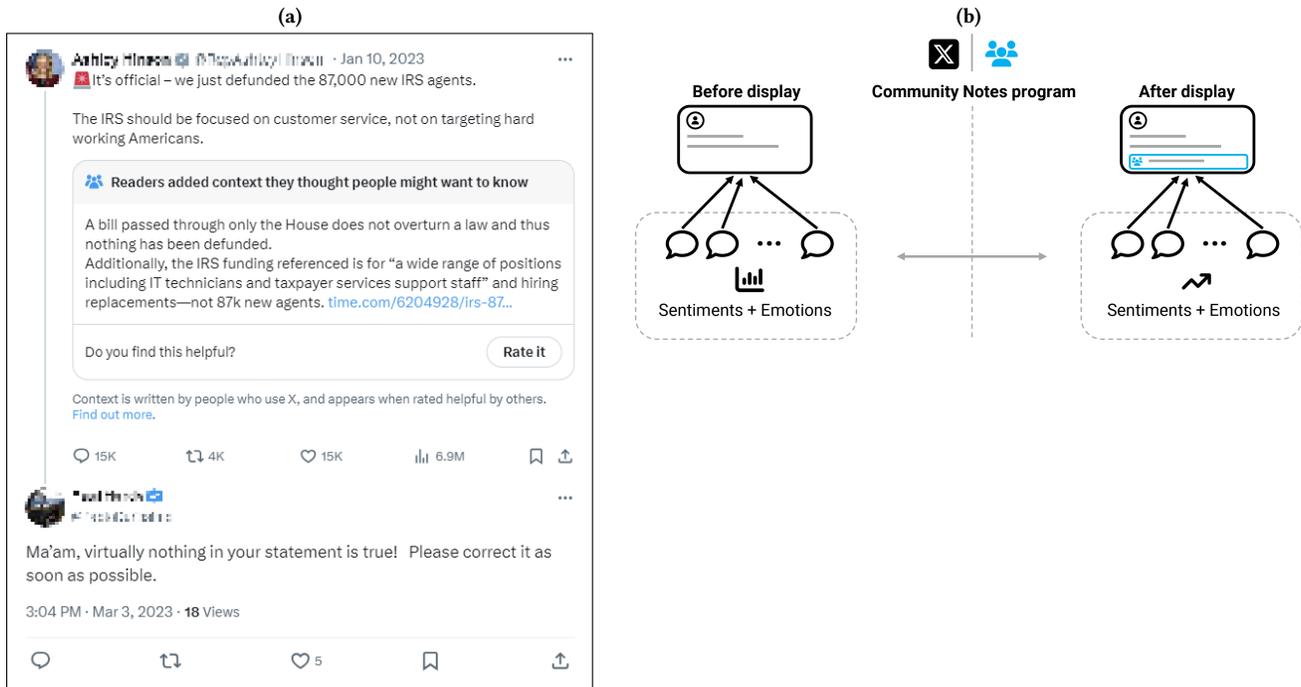


Figure 1: Research overview. (a) An example of a misleading post with a displayed community note and one direct reply after note display. (b) Illustration of our research setup. Before note display, users who reply to the source posts can only see the original post content. After note display, users who reply to the source posts can see both original post content and displayed community notes flagging the post content as misleading. Using the sentiments and emotions in replies before note display as a baseline, we can examine the changes in sentiments and emotions in replies after the display of community notes.

3.1.3 Reply dataset. We retrieve the entire conversation thread for each source post via X’s Academic Research API.³ Each conversation thread contains all the replies to the original posts, and they share one unique conversation ID. In this study, we analyze the changes of sentiments and emotions in replies to the original misleading posts after the display of community notes. Therefore, we only consider direct replies to the source posts and omit indirect replies (i. e., replies to replies). We successfully pull $N = 2,225,260$ direct replies to 1841 misleading source posts (94.3%) that were fact-checked via community notes during a period of 4 months⁴ since its roll-out (Table 1). The remaining 5.7% (112) misleading source posts have no (English) replies and are thus omitted.

3.2 Sentiments and Basic Emotions

We extract sentiment and emotion features in the source posts and their replies using two state-of-the-art machine learning models. Specifically, we compute the probabilities of positive and negative sentiments for each source post and reply using the Twitter-roBERTa-base model (2022 updated). It was trained on 124 million posts created from January 2018 to December 2021, and fine-tuned

³See the documentation for X’s API at <https://developer.x.com/en/docs/twitter-api/conversation-id>.

⁴The collection of our dataset was completed over two months, from April to June 2023. Subsequently, X deprecated the Academic Research API, restricting our ability to access additional data for analysis. Notwithstanding, previous research implies that our dataset covering the initial four-month period provides a representative data basis for the Community Notes program [19, 20, 97].

for sentiment analysis.⁵ This model achieves a high prediction performance compared to other sentiment models [16, 73]. Subsequently, we compute the probabilities of Ekman’s six basic emotions (i. e., anger, disgust, fear, joy, sadness, and surprise) [40, 106] for each source post and reply using a distilled version of the RoBERTa-base model.⁶ The model was pre-trained on OpenWebTextCorpus, a reproduction of OpenAI’s WebText dataset with more than 8 million documents, and fine-tuned for emotion analysis [56]. The performance of this emotion analysis model is comparable to human annotations and has been used in several previous studies [14, 102]. Summary statistics for the sentiments and emotions in the source posts and replies are reported in Table 1.

3.3 Political vs. Non-Political Posts

Misleading posts in the political domain are particularly concerning due to their potential to undermine democratic processes and destabilize societies [37]. As observed by previous studies [77, 86, 121], political misinformation spreads faster than non-political misinformation and is particularly resistant to corrections. Therefore, we additionally conduct a sensitivity analysis examining changes in sentiments and emotions in replies to misleading posts that cover

⁵The sentiment model is available at <https://huggingface.co/cardiffnlp/twitter-roberta-base-sentiment-latest>.

⁶The emotion model is available at <https://huggingface.co/j-hartmann/emotion-english-distilroberta-base>.

Table 1: Dataset overview. Reported are mean values or count numbers for the variables (standard deviations in parentheses).

	(1) All	(2) Political	(3) Non-political
#Replies	2,225,260	1,068,732	1,156,528
#Source posts	1841	477	1364
Date range of replies	2022-12-11 – 2023-05-28	2022-12-11 – 2023-05-27	2022-12-11 – 2023-05-28
Date range of source posts	2022-12-11 – 2023-04-05	2022-12-11 – 2023-04-05	2022-12-11 – 2023-04-05
<u>Sentiments in replies</u>			
Positive	0.112 (0.222)	0.098 (0.209)	0.124 (0.233)
Negative	0.565 (0.336)	0.608 (0.324)	0.525 (0.342)
<u>Sentiments in source posts</u>			
Source positive	0.114 (0.222)	0.084 (0.184)	0.125 (0.233)
Source negative	0.443 (0.334)	0.513 (0.315)	0.419 (0.337)
<u>Emotions in replies</u>			
Anger	0.181 (0.260)	0.205 (0.272)	0.160 (0.246)
Disgust	0.097 (0.183)	0.098 (0.181)	0.096 (0.185)
Fear	0.050 (0.146)	0.047 (0.138)	0.054 (0.153)
Joy	0.066 (0.170)	0.061 (0.163)	0.069 (0.176)
Sadness	0.107 (0.192)	0.109 (0.193)	0.104 (0.192)
Surprise	0.172 (0.247)	0.164 (0.241)	0.180 (0.252)
<u>Emotions in source posts</u>			
Source anger	0.159 (0.243)	0.195 (0.260)	0.147 (0.236)
Source disgust	0.041 (0.117)	0.041 (0.112)	0.040 (0.119)
Source fear	0.173 (0.276)	0.151 (0.242)	0.181 (0.287)
Source joy	0.077 (0.176)	0.068 (0.158)	0.080 (0.182)
Source sadness	0.126 (0.211)	0.140 (0.216)	0.122 (0.209)
Source surprise	0.175 (0.245)	0.174 (0.249)	0.175 (0.244)

a political vs. a non-political topic. To do so, we extract topic domains from the context annotations attached to each source post and classify them as political or non-political. X performs entity recognition and semantic analysis (e. g., keywords, hashtags, and handles) for all posts and assigns corresponding domain labels based on a topic domain list.⁷ Within the domain list, we consider “Politician,” “Political Body,” and “Political Race” as political topic domains. Posts containing one of the political topic domains are classified as political; otherwise, they are classified as non-political. As a result, there are 477 political misleading posts in our data, accounting for 25.9% of all misleading posts. However, nearly half of the replies are directed at political misleading posts (see Table 1).

3.4 Empirical Models

In this study, we longitudinally examine the direct replies to the source posts before and after the display of community notes. To address our research questions, we use Regression Discontinuity Design (RDD) to causally estimate the effects of the display of community notes in terms of sentiments and basic emotions. RDD is a quasi-experimental method to estimate the average treatment effect of an intervention by exploiting a discontinuity in the relationship between a running variable and an outcome variable at a

specific threshold or cutoff point [5, 17, 21, 61]. In the following, we introduce the dependent variables (Section 3.4.1), display indicator and running variable (Section 3.4.2), and additional independent variables (Section 3.4.3) in our analysis. Subsequently, we detail the specification of our empirical RDD models (Section 3.4.4).

3.4.1 Dependent variables. We define the following dependent variables that measure expressions of sentiments, basic emotions, and moral outrage in replies to social media posts:

- Dependent variables for positive and negative sentiments
 - **Positive:** A continuous variable from 0 to 1 that indicates the probability of positive sentiment in a reply.
 - **Negative:** A continuous variable from 0 to 1 that indicates the probability of negative sentiment in a reply.
- Dependent variables for basic emotions
 - **Anger:** A continuous variable from 0 to 1 that indicates the probability of anger in a reply.
 - **Disgust:** A continuous variable from 0 to 1 that indicates the probability of disgust in a reply.
 - **Fear:** A continuous variable from 0 to 1 that indicates the probability of fear in a reply.
 - **Joy:** A continuous variable from 0 to 1 that indicates the probability of joy in a reply.
 - **Sadness:** A continuous variable from 0 to 1 that indicates the probability of sadness in a reply.
 - **Surprise:** A continuous variable from 0 to 1 that indicates the probability of surprise in a reply.

⁷The topic domains are curated by domain experts and have undergone refinement over several years. X’s internal team audits this approach quarterly through the manual review of 10,000 posts across all the domains. The precision scores are consistently ~80%. The topic domain list is available at <https://developer.x.com/en/docs/twitter-api/annotations/overview>.

- Dependent variable for moral outrage
 - **MoralOutrage**: A continuous variable from 0 to 1 that indicates the probability of moral outrage in a reply. Moral outrage represents a unique combination of anger and disgust in which the two emotions interact so that each exacerbates the other’s influence on judgments [104]. This interaction extends beyond the isolated effects of anger and disgust, thereby offering important theoretical insights into moral outrage [9, 28, 104]. Accordingly, we use the product of anger and disgust (i. e., $Anger \times Disgust$) as a measure of moral outrage to explore their interactive effect.⁸

3.4.2 Display indicator and running variable. Based on RDD, we aim to examine whether there are significant discontinued changes in sentiments and emotions after the display of community notes and estimate the effects. The key rationale behind this is that the sentiments and emotions in replies should change smoothly (or keep stable) over time if without the display of community notes. Given this, we define the display indicator of community notes (binary treatment indicator) as **Displayed**. It is equal to 1 if a reply is created after note display; otherwise, it is equal to 0. Additionally, we examine the changes in sentiments and emotions over time before and after the display of community notes. Therefore, we define the running variable as **HoursFromDisplay**, which is a continuous variable indicating the hour(s) from note display to reply creation.

3.4.3 Additional independent variables. Given that connected social media users may share similar emotions [51, 69, 100], the sentiments and emotions in the source posts are likely to affect the sentiments and emotions in the replies from users who are exposed to these source posts (see correlation analysis in Suppl. S1.1). Thus, we incorporate source sentiments and emotions into the regression models as independent variables:

- Independent variables for source sentiments in misleading posts
 - **SourcePositive**: A continuous variable from 0 to 1 that indicates the probability of positive sentiment in a source post.
 - **SourceNegative**: A continuous variable from 0 to 1 that indicates the probability of negative sentiment in a source post.
- Independent variables for source emotions in misleading posts
 - **SourceAnger**: A continuous variable from 0 to 1 that indicates the probability of anger in a source post.
 - **SourceDisgust**: A continuous variable from 0 to 1 that indicates the probability of disgust in a source post.
 - **SourceFear**: A continuous variable from 0 to 1 that indicates the probability of fear in a source post.
 - **SourceJoy**: A continuous variable from 0 to 1 that indicates the probability of joy in a source post.
 - **SourceSadness**: A continuous variable from 0 to 1 that indicates the probability of sadness in a source post.
 - **SourceSurprise**: A continuous variable from 0 to 1 that indicates the probability of surprise in a source post.

In addition, the times when community notes are displayed on the source post are not fixed and depend on note helpfulness scores that are recalculated every hour. Therefore, changes in sentiments and emotions after note display might be confounded by the post age

⁸As a robustness check, we repeated our analysis with a dedicated classifier for digital outrage developed by previous research [9]. Here, we observe quantitatively identical results (see details in Suppl. S6).

from post creation to reply creation. Given this, we use **PostAge** to control for the possible confounding effects of post age. Descriptive statistics of sentiments and emotions in replies and source posts are summarized in Table 1.

3.4.4 Regression models. We specify two RDD models to estimate the effects of the display of community notes on sentiments and emotions in replies, respectively. For positive and negative sentiments, the RDD model is specified as:

$$\begin{aligned} Sentiment_i = & \beta_0 + \beta_1 Displayed_i + \beta_2 HoursFromDisplay_i \\ & + \beta_3 SourcePositive_i + \beta_4 SourceNegative_i \\ & + \beta_5 PostAge_i + \epsilon_i, \end{aligned} \quad (1)$$

where *Sentiment* denotes the dependent variable of *Positive* or *Negative*, β_0 is the intercept, and ϵ is the residual. β_1 to β_5 are coefficient estimates for the independent variables.

For the six basic emotions, the RDD model is specified as:

$$\begin{aligned} Emotion_i = & \beta_0 + \beta_1 Displayed_i + \beta_2 HoursFromDisplay_i \\ & + \beta_3 SourceAnger_i + \beta_4 SourceDisgust_i \\ & + \beta_5 SourceFear_i + \beta_6 SourceJoy_i \\ & + \beta_7 SourceSadness_i + \beta_8 SourceSurprise_i \\ & + \beta_9 PostAge_i + \epsilon_i, \end{aligned} \quad (2)$$

where *Emotion* denotes the dependent variable of one of the six basic emotions: *Anger*, *Disgust*, *Fear*, *Joy*, *Sadness*, or *Surprise*. Additionally, we further take *MoralOutrage* (i. e., $Anger \times Disgust$) as the dependent variable to explore the changes in moral outrage after the display of community notes. The independent variables for source sentiments are replaced with the variables for source emotions.

To balance the scales of variables, we further z-standardize all continuous variables when fitting the RDD models. The standardization is based on:

$$z_i = \frac{x_i - \mu}{\sigma}, \quad (3)$$

where x_i denotes original value for a continuous variable X , and z_i is the standardized value. μ and σ are mean and standard deviation of the variable X , respectively. Therefore, the coefficient estimates for *Displayed*, i. e., β_1 , in the models cannot be directly interpreted as the effects of note display. Given this, we define that Y_{before} is the outcome of a dependent variable before the display of community notes and Y_{after} is the outcome of a dependent variable after the display of community notes. Assuming that all the other independent variables are equal for the outcomes of Y_{before} and Y_{after} , we have:

$$\begin{aligned} Y_{after} - Y_{before} &= \beta_1 \\ \Rightarrow \frac{Y_{after} - \mu}{\sigma} - \frac{Y_{before} - \mu}{\sigma} &= \beta_1 \\ \Rightarrow Y_{after} - Y_{before} &= \sigma \beta_1. \end{aligned} \quad (4)$$

The value of $\sigma \beta_1$ can represent the estimated (predicted) effects of the display of community notes on the sentiments and emotions, with all else being equal.

Considerations for the adoption of RDD: Previous studies evaluating the effect of community notes on *engagement* (e. g., the number of reposts) have typically employed Difference-in-Differences (DiD) methods by constructing a counterfactual control group that evolves in parallel with the treatment group prior to

the display of community notes [19–21, 97]. However, constructing such control groups often omits a significant portion of posts. While our dataset includes a large volume of replies, the limited number of source posts makes it infeasible to conduct such matching at the post level. Additionally, previous studies focusing on engagement typically analyze repost counts at an hourly level to examine engagement reductions attributed to the display of community notes. Therefore, they do not require the collection of every individual reply with a precise timestamp, and can count the number of reposts as zero for intervals without repost activity. In contrast, this paper focuses on analyzing sentiments and emotions in individual replies to misleading posts before and after the display of community notes. This analysis necessitates collecting every reply along with its associated timestamp and content. As sentiments and emotions cannot be assumed to be zero in intervals without replies, gaps in time units are likely to arise. These considerations lead us to adopt RDD with a continuous running variable in time [17, 61].

Further, research has discussed the use of time as a running variable in RDD, i. e., Regression Discontinuity in Time (RDiT), and suggested considering the time-series nature of the underlying data-generating process [58, 66]. On social media platforms, the autoregressive time-series structure originates from the creation of original posts, extends to direct replies, and further encompasses indirect replies (i. e., replies to replies), with each direct reply initiating a separate reply thread. To minimize complications from autoregressive processes within reply threads, our study focuses exclusively on direct replies to the original posts and omits indirect replies. Nevertheless, direct replies to the same source post across different observational units are still likely to exhibit correlations in terms of sentiments and emotions. However, according to X’s algorithmic rules [125], replies in the conversations associated with the same source post are typically not displayed in chronological order. Thus, the correlations among direct replies are likely to primarily arise from shared factors, e. g., source sentiments and emotions, in the original posts (as illustrated in Fig. 1b).⁹ Following prior research [58, 121], we cluster standard errors in all regression models by unique source posts to address the potential violation of independence assumption. By clustering the standard errors, we account for the potential within-cluster correlations and provide more precise and reliable standard error estimates. In addition to the cluster-robust method, the incorporation of source sentiments and emotions in the regression models explicitly controls for the correlations among replies linked to the same source posts.

4 Results

4.1 Data Overview

Our dataset covers four months since the roll-out of the “Community Notes” program on X in December of 2022. Fig. 2a shows that the daily number of misleading posts on which community notes are displayed has an increasing trend, with an average of 15.7. Fig. 2b shows that misleading source posts tend to have higher negative

sentiment (mean of 0.443), compared to positive sentiment (mean of 0.114; $t = -28.810$, $p < 0.001$). In terms of emotions, anger (mean of 0.159), fear (mean of 0.173), and surprise (mean of 0.175) are most common in the fact-checked misleading source posts (Fig. 2c).

To examine the effects of community notes on the sentiments and emotions in replies, we collect a large-scale longitudinal dataset of replies to misleading posts. Fig. 2d shows that each misleading source post receives an average of 1209 (median of 298) replies in total (i. e., before and after the display of community notes). On average, community notes are displayed 52.5 hours (median of 24.6 hours) after the creation of the source misleading posts (Fig. 2e). Notably, most of the replies (mean of 82.2%, median of 88.8%) are created before the display of community notes (Fig. 2f). Therefore, some misleading posts may not have sufficient reply samples after note display to perform comparisons with the replies before note display. Given this, we restrict our main analysis to misleading posts that have at least (\geq) 5 replies both before and after the display of community notes. This criterion¹⁰ results in 1339 misleading posts being included in the comparison of sentiments and emotions in replies before and after note display.

The key assumption of RDD in this paper is that a discontinuity exists around the cut-off point, i. e., the time when community notes are displayed [17, 61, 109]. In the absence of this intervention, the sentiments and emotions in replies are expected to remain continuous or stable. To validate this assumption, we first examine changes in sentiments and emotions over time before the display of community notes. Our analysis shows that, in the absence of community notes, the sentiments and emotions in replies remain relatively stable and exhibit no statistically significant relationship with the age of misleading posts (see details in Suppl. S1.2). Subsequently, we use an interrupted time series design to analyze potential discontinuities in sentiments and emotions in replies following the display of community notes. The results reveal statistically significant discontinuities for negative sentiment, anger, and surprise in replies after the display of community notes (each $p < 0.05$). Additionally, aside from the identified discontinuity at the cut-off point, the sentiments and emotions in replies after note display show trends similar to those in replies before note display and remain stable over time (see details in Suppl. S1.2). These findings confirm that RDD is appropriate for examining the effects of community notes on sentiments and emotions in replies.

Another important consideration for RDD is to include as many samples as possible within a relatively short bandwidth, aiming to mitigate potential confounding factors associated with larger windows around the display of community notes [17, 58, 61]. In our dataset, 97.8% of total replies are captured within a one-week window before and after the display of community notes. We consider 10% of this window (i. e., 16 hours) as the bandwidth, in which 61.3% of total replies are included, to examine the effect sizes of community notes on sentiments and emotions in our main analysis.¹¹ To further illustrate the discontinuity around the display of community notes, we plot the hourly averages of sentiments and

⁹In our later robustness checks, we validate that the correlations of sentiments and emotions among direct replies are stable and consistent across different temporal orders. The results confirm that the correlations arise from the influence of original posts rather than from temporal dependencies between direct replies. Furthermore, we repeated our analysis with lagged-dependent covariates as an additional check. Again, the results remain robust and support our findings (see details in Suppl. S7.2).

¹⁰This criterion is primarily applied to improve statistical inference. Our results remain robust and consistent when this restriction is removed (see details in Suppl. S7.4).

¹¹Given the stable trends of sentiments and emotions in replies over time, we verify the robustness of RDD estimates across different bandwidths and find qualitatively identical results throughout the entire reply lifespan (see Fig. 5).

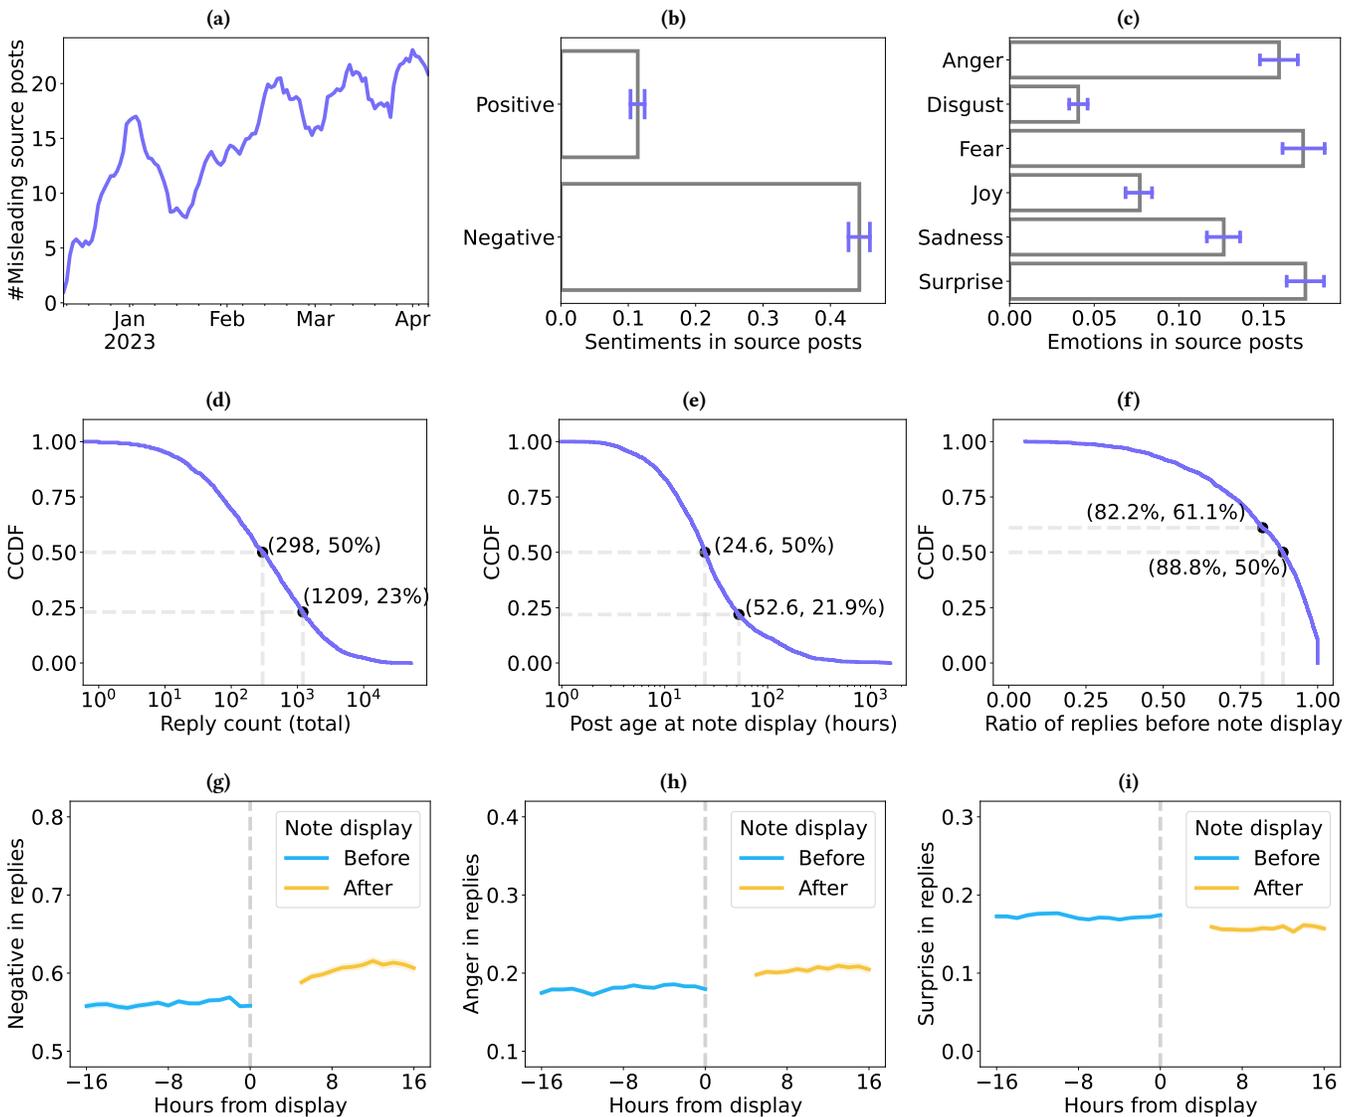


Figure 2: Summary statistics for misleading source posts, direct replies, and display of community notes. (a) The 2-week rolling average daily number of misleading source posts that are attached with displayed community notes. (b) The means of positive and negative sentiments in the misleading source posts. (c) The means of anger, disgust, fear, joy, sadness, and surprise in the misleading source posts. (d) The Complementary Cumulative Distribution Function (CCDF) for the number of replies that are directed at each misleading source post. (e) The CCDF for hour(s) from post creation to note display. (f) The CCDF for the ratio of replies before note display to the total replies. (g) The hourly averages of negative sentiment in replies across hours from note display. (h) The hourly averages of anger in replies across hours from note display. (i) The hourly averages of surprise in replies across hours from note display. The hourly averages of positive sentiment and other emotions are shown in Fig. S2 in Suppl. S1.2. The error bars (bands) represent 95% Confidence Intervals (CIs). Notably, previous studies consistently suggest a potential cold-start period of 4 hours for community notes to reach their full effect [19, 97]. Therefore, we visually omit reply points within the initial four hours after the display of community notes for better readability.

emotions in replies within the 16-hour bandwidth before and after note display. Figs. 2g–2i demonstrate the significant discontinuities in negative sentiment, anger, and surprise at the cut-off point with stable levels at both sides (see Fig. S2 in Suppl. S1.2 for the visualizations of positive sentiment and other emotions).

To shed light on the relationships between sentiments and basic emotions in our reply dataset, we examine their cross-correlations based on Pearson’s r correlation coefficients. We find that anger ($r = 0.406$, $p < 0.001$), disgust ($r = 0.323$, $p < 0.001$), fear ($r = 0.089$, $p < 0.001$), and sadness ($r = 0.192$, $p < 0.001$) have significantly

positive correlations with negative sentiment, while joy ($r = 0.572$, $p < 0.001$) and surprise ($r = 0.044$, $p < 0.001$) are positively correlated with positive sentiment (see Table S5 in Suppl. S1 for all cross-correlations). This suggests that, in replies to misleading posts, negative feelings are often expressed through anger, disgust, fear, and sadness. Notably, anger and disgust have the strongest correlations with negative sentiment. We further examine the cross-correlations among the six basic emotions in each reply. We find that only anger and disgust have a significantly positive correlation with each other ($r = 0.111$, $p < 0.001$). Other emotions all show negative correlations (reported in Table S6 in Suppl. S1). This suggests that replies to misleading posts often express anger and disgust together, which is a strong indicator of moral outrage [104].

4.2 Analysis of Sentiments and Basic Emotions (RQ1.1)

4.2.1 Descriptive analysis. We start by using t -tests to descriptively report changes in sentiments and emotions in replies after the display of community notes, compared to before note display.¹² For sentiments, positive sentiment in replies after the display of community notes (mean of 0.115) is significantly lower than before the note display (mean of 0.127; $t = 10.249$, $p < 0.001$; Fig. S3a). In contrast, negative sentiment in replies after the display of community notes (mean of 0.566) is significantly higher than before the note display (mean of 0.524; $t = -23.748$, $p < 0.001$; Fig. S3b). In terms of basic emotions, anger (mean of 0.179) and disgust (mean of 0.093) in replies after the display of community notes are significantly higher than anger (mean of 0.155; $t = -18.531$, $p < 0.001$; Fig. S3c) and disgust (mean of 0.090; $t = -2.717$, $p < 0.01$; Fig. S3d) in replies before the note display, respectively. Fear have no statistically significant changes after the display of community notes (each $p > 0.05$; see Figs. S3e–S3g). Surprise in replies after the display of community notes (mean of 0.179) is significantly lower than that in replies before the note display (mean of 0.185; $t = 5.325$, $p < 0.001$; Fig. S3h). Taken together, our descriptive analysis suggests that replies after note display tend to be more negative and embed higher anger and disgust, compared to replies before the display of community notes (see details in Suppl. S2).

4.2.2 Causal analysis. We now use regression models based on RDD to estimate the causal changes in sentiments and emotions in replies to source misleading posts that are attributed to the display of community notes. The coefficient estimates for the regressions across sentiments and basic emotions are visualized in Fig. 3. Additionally, to interpret the effect sizes of the display of community notes on the specific sentiments and emotions in replies to misleading posts, we examine the extent to which the sentiments and basic emotions in replies change compared to their corresponding baselines before the display of community notes. The changes in sentiments and emotions in replies to misleading posts within the 16-hour window, along with the predicted effects of displaying community notes, are visualized in Fig. 4.

¹²We use the t -test rather than non-parametric alternatives such as the Mann–Whitney U -test or the Kolmogorov–Smirnov test due to its advantage in reporting confidence intervals, which enhances the robustness of statistical testing for large-scale data. Although the distributions of emotions and sentiments are not strictly normal, their sample means are approximately normally distributed in our dataset. Therefore, the t -test is valid in our study based on the central limit theorem [44, 74].

Effect of community notes display: For positive sentiment in replies, the coefficient estimate of *Displayed* is significantly negative (Fig. 3a: *coef.* = -0.055 , $p < 0.01$; 95% CI: $[-0.089, -0.021]$). This means that the display of community notes decreases the positive sentiment in replies by 0.012 (95% CI: $[-0.020, -0.005]$). As shown in Fig. 4a, the baseline of positive sentiment in replies before note display is 0.125, and the positive sentiment in replies after note display is predicted to be 0.113 (95% CI: $[0.106, 0.121]$). This suggests that positive sentiment in replies decreases by 9.7% after the display of community notes compared to that before note display. In contrast, for negative sentiment in replies, the coefficient estimate of *Displayed* is significantly positive (Fig. 3b: *coef.* = 0.114 , $p < 0.001$; 95% CI: $[0.070, 0.158]$). This means that the display of community notes increases the negative sentiment in replies by 0.038 (95% CI: $[0.023, 0.053]$). As shown in Fig. 4b, the baseline of negative sentiment in replies before note display is 0.521, and the negative sentiment in replies after note display is predicted to be 0.559, (95% CI: $[0.545, 0.574]$). This suggests that negative sentiment in replies increases by 7.3% after the display of community notes compared to that before note display.

We further examine the effects of community notes display on changes in basic emotions in replies to misleading posts before and after the display of community notes. The results are as follows:

- For anger in replies, the coefficient estimate of *Displayed* is significantly positive (Fig. 3c: *coef.* = 0.078 , $p < 0.001$; 95% CI: $[0.058, 0.098]$). This means that the display of community notes increases anger in replies by 0.020 (95% CI: $[0.015, 0.026]$). As shown in Fig. 4c, the baseline of anger in replies before note display is 0.154, and the predicted anger in replies after note display is 0.174 (95% CI: $[0.169, 0.180]$). This suggests that anger in replies increases by 13.2% after the display of community notes compared to that before note display.
- For disgust in replies, the coefficient estimate of *Displayed* is significantly positive (Fig. 3d: *coef.* = 0.023 , $p < 0.05$; 95% CI: $[0.004, 0.042]$). This means that the display of community notes increases disgust in replies by 0.004 (95% CI: $[0.001, 0.008]$). As shown in Fig. 4d, the baseline of disgust in replies before note display is 0.090, and the predicted disgust in replies after note display is 0.094 (95% CI: $[0.091, 0.098]$). This suggests that disgust in replies increases by 4.7% after the display of community notes compared to that before note display.
- For fear, joy, and sadness in replies, the coefficient estimates of *Displayed* are all not statistically significant (each $p > 0.05$; see Figs. 3e–3g). This means that the display of community notes has no significant effect on these three emotions in replies to misleading posts.
- For surprise in replies, the coefficient estimate of *Displayed* is significantly negative (Fig. 3h: *coef.* = -0.043 , $p < 0.001$; 95% CI: $[-0.060, -0.025]$). This means that the display of community notes decreases surprise in replies by 0.011 (95% CI: $[-0.015, -0.006]$). As shown in Fig. 4h, the baseline of surprise in replies before note display is 0.184, and the predicted surprise in replies after note display is 0.174 (95% CI: $[0.169, 0.178]$). This suggests that surprise in replies decreases by 5.7% after the display of community notes compared to that before note display.

Taken together, the display of community notes increases 7.3% more negative sentiment and triggers 13.2% more anger and 4.7% more

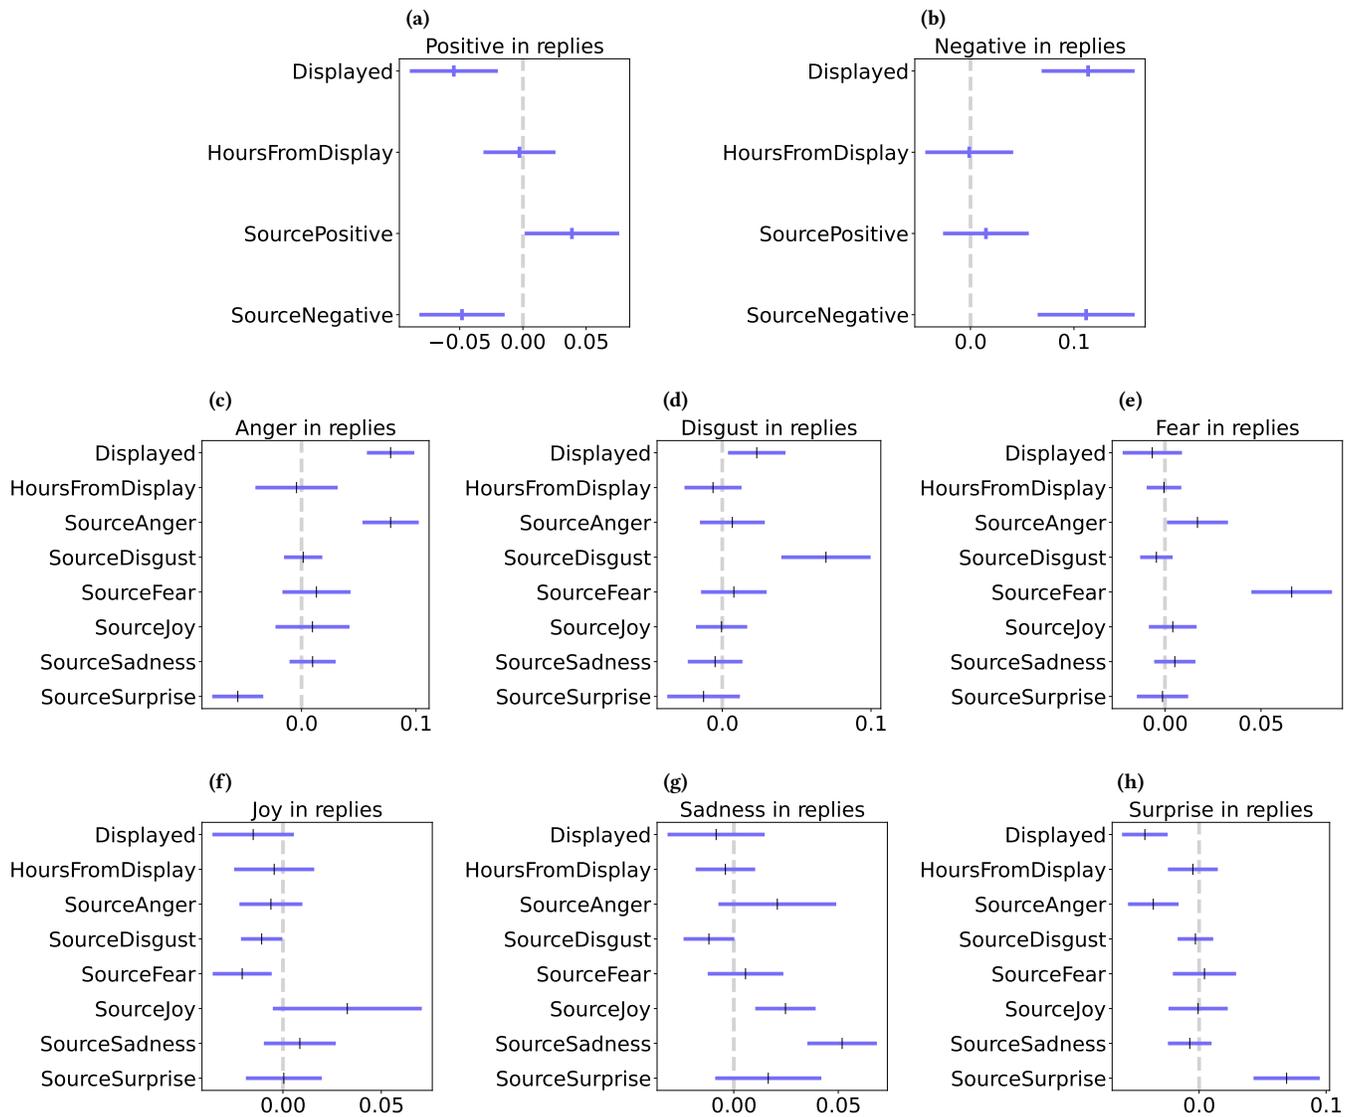


Figure 3: The estimated coefficients for the independent variables – *Displayed*, *HoursFromDisplay*, and source sentiments (or emotions). The independent variable *PostAge* is included during estimation but omitted in the visualization for better readability. Shown are mean values with error bars representing 95% CIs. Standard errors are clustered by source posts. The dependent variables are (a) positive sentiment in replies, (b) negative sentiment in replies, (c) anger in replies, (d) disgust in replies, (e) fear in replies, (f) joy in replies, (g) sadness in replies, and (h) surprise in replies, respectively. The full estimation results are reported in Suppl. S3.

disgust in replies to misleading posts, respectively. The effects of community notes display on sentiments and emotions in replies are concordant with the observed changes of sentiments and emotions in replies after the display of community notes (see Fig. S3 in Suppl. S2).

Effect of sentiments and emotions in source posts: We further examine the link between the original sentiments and emotions in the source misleading posts and the subsequent sentiments and emotions in the direct replies. For positive sentiment in replies, the

coefficient estimate of *SourcePositive* in Fig. 3a is significantly positive ($coef. = 0.039, p < 0.05; 95\% CI: [0.002, 0.075]$); for negative sentiment in replies, the coefficient estimate of *SourceNegative* in Fig. 3b is also significantly positive ($coef. = 0.112, p < 0.001; 95\% CI: [0.066, 0.158]$). This means that the positive and negative sentiments in replies are positively linked to the original positive and negative sentiments in the source misleading posts, respectively. In terms of emotions, the coefficient estimate of *SourceAnger* in Fig. 3c ($coef. = 0.078, p < 0.001; 95\% CI: [0.054, 0.102]$), the coefficient estimate of *SourceDisgust* in Fig. 3d ($coef. = 0.070, p < 0.001; 95\%$

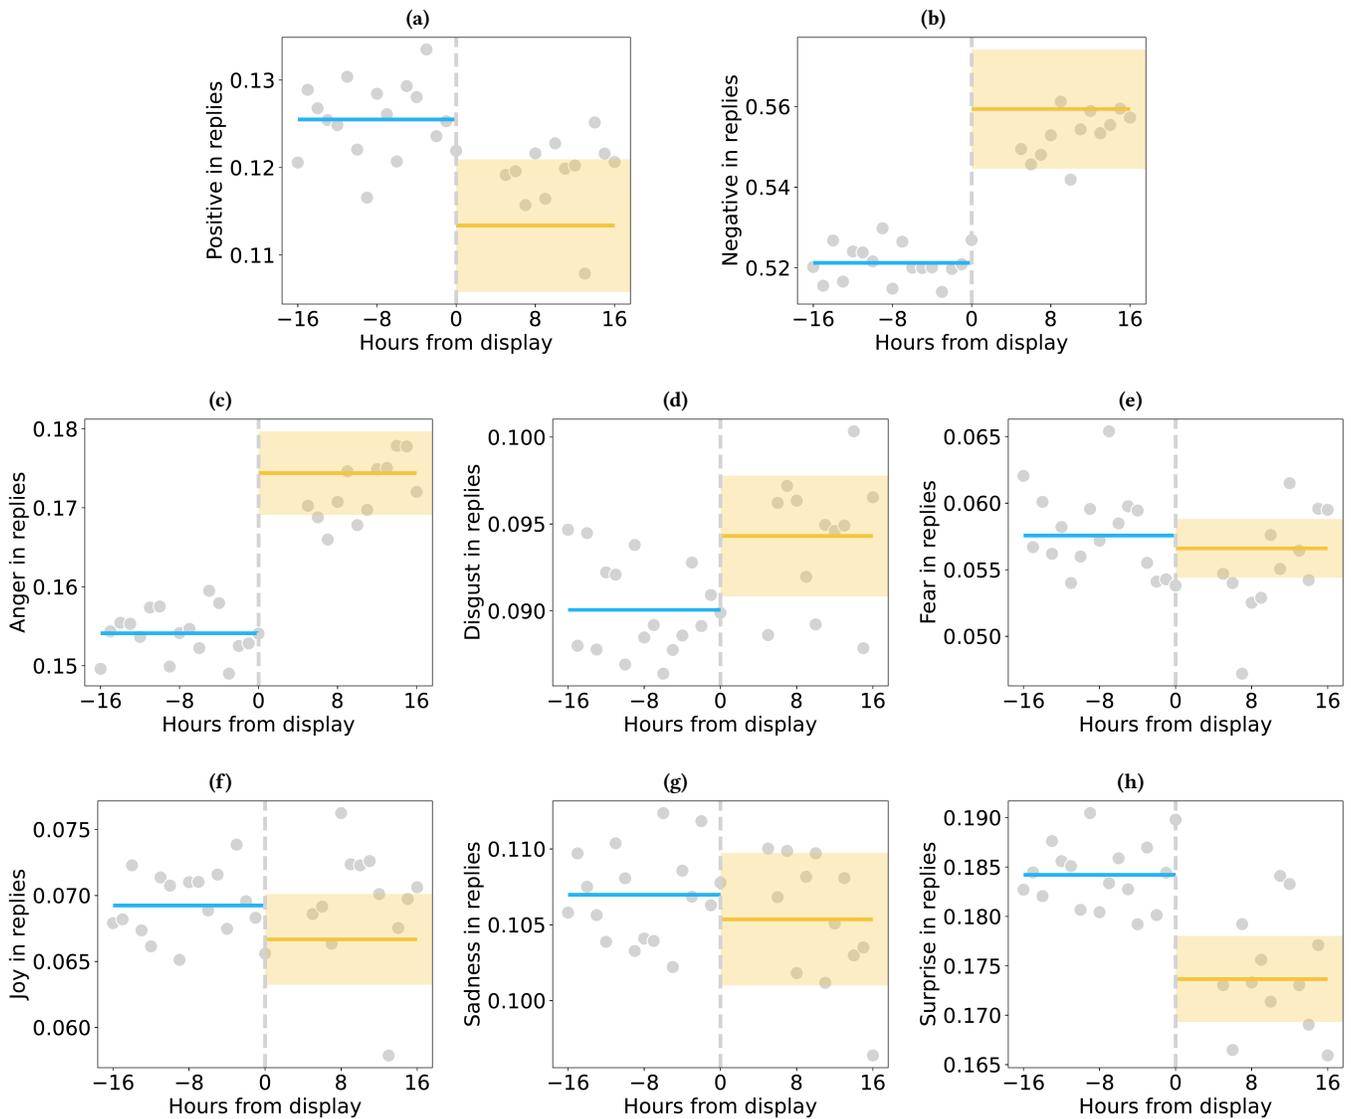


Figure 4: The predicted effects of the display of community notes on sentiments and emotions in replies. (a) The predicted effect on positive sentiment in replies. (b) The predicted effect on negative sentiment in replies. (c) The predicted effect on anger in replies. (d) The predicted effect on disgust in replies. (e) The predicted effect on fear in replies. (f) The predicted effect on joy in replies. (g) The predicted effect on sadness in replies. (h) The predicted effect on surprise in replies. The grey points indicate the hourly averages of sentiments or emotions between 16 hours before and 16 hours after the display of community notes. Similar to Fig. 2, we visually omit reply points within the initial four hours after the display of community notes for better readability. The blue lines indicate the averages of sentiments or emotions in replies over the 16 hours before note display and represent the baselines during before-display period. The yellow lines indicate the averages of sentiments or emotions over the 16 hours after the display of community notes. The average at the yellow line is the sum of the predicted effect and the corresponding baseline at the blue line in each figure. The yellow bands represent 95% CIs. The predicted effects are transformed based on the coefficient estimates of *Displayed* (see Fig. 3). The full estimation results are reported in Suppl. S3.

CI: [0.040, 0.099]), the coefficient estimate of *SourceFear* in Fig. 3e (*coef.* = 0.066, $p < 0.001$; 95% CI: [0.045, 0.087]), the coefficient estimate of *SourceSadness* in Fig. 3g (*coef.* = 0.052, $p < 0.001$; 95% CI: [0.036, 0.068]), and the coefficient estimate of *SourceSurprise* in Fig. 3h (*coef.* = 0.069, $p < 0.001$; 95% CI: [0.043, 0.094]) are all

significantly positive. This suggests that the expressions of anger, disgust, fear, sadness, and surprise in source misleading posts have significantly positive effects on anger, disgust, fear, sadness, and surprise in the subsequent replies, respectively. However, the coefficient estimate of *SourceJoy* in Fig. 3f is not statistically significant

($coef. = 0.033$, $p = 0.084$; 95% CI: $[-0.004, 0.070]$). Thus, we find no evidence that the expression of joy in replies is significantly linked to the expression of joy in the source misleading posts.

Effect of reply timing: We examine whether the sentiments and emotions in replies change over time without the display of community notes. We find that the coefficient estimates of *HoursFromDisplay* are not statistically significant across all dependent variables (each $p > 0.05$; see Figs. 3a–3h). The coefficient estimates of *PostAge* are also not statistically significant (see details in Suppl. S3). This means that the sentiments and emotions in replies are relatively stable over time if community notes are not displayed on the corresponding misleading posts.

RDD estimates across different bandwidths: Based on the previous analysis of stable trends in sentiments and emotions within the reply lifespan (see Section 4.1), the RDD estimates across different bandwidths should remain robust and have no statistically significant differences. To validate this, we examine RDD estimates using three different bandwidths: 16 hours (16H), one week (1W), and the entire reply lifespan (LSP). As shown in Fig. 5, the RDD coefficient estimates across sentiments and emotions in replies have no statistically significant differences within the three bandwidths (see full estimation results in Suppl. S4). Additionally, we conduct *SUEST* tests to compare RDD models across the three bandwidths, further confirming the equivalence of the common coefficient estimates within these models (each $p > 0.05$, see Table 2).

Summary of findings: In sum, we find that replies are more negative and express more anger and disgust after the display of community notes compared to those before the display. Specifically, the display of community notes triggers 13.2% more anger and 4.7% more disgust in replies to source misleading posts. Additionally, the sentiments and almost all emotions (except for joy) in replies are positively linked to those expressed in the source misleading posts. Furthermore, we find that the expressions of sentiments and emotions in replies are relatively stable over time if without the display of community notes.

4.3 Analysis Across Politic and Non-Political Misleading Posts (RQ1.2)

We conduct sensitivity analysis to examine whether the effects of community notes on sentiments and emotions in replies vary across political and non-political misleading posts. The coefficient estimates for the sensitivity analysis across political and non-political misleading posts are reported in Fig. 6.

We first examine the sensitivity of the effects of community notes on positive and negative sentiments. In Fig. 6a, the coefficient estimate of *Displayed* for positive sentiment in replies to political misleading posts is significant and negative ($coef. = -0.068$, $p < 0.01$; 95% CI: $[-0.115, -0.022]$). However, the coefficient estimate of *Displayed* for positive sentiment in replies to non-political misleading posts is not statistically significant ($coef. = -0.020$, $p = 0.197$; 95% CI: $[-0.051, 0.010]$). This means that the negative effect of community notes on positive sentiment in replies is limited to political misleading posts. Regarding the negative sentiment in replies (Fig. 6b), the coefficient estimates of *Displayed* are significantly positive in both political ($coef. = 0.116$, $p < 0.001$; 95% CI: $[0.067, 0.164]$) and non-political misleading posts ($coef. = 0.075$,

$p < 0.001$; 95% CI: $[0.042, 0.107]$). This means that the effects of community notes generally apply to both political and non-political misleading posts.

Next, we examine the sensitivity of the effects of community notes on basic emotions:

- In terms of anger in replies (Fig. 6c), the coefficient estimates of *Displayed* are significantly positive in both political ($coef. = 0.063$, $p < 0.001$; 95% CI: $[0.036, 0.091]$) and non-political misleading posts ($coef. = 0.070$, $p < 0.001$; 95% CI: $[0.042, 0.099]$). This suggests that the display of community notes triggers more anger in replies to both political and non-political misleading posts.
- In Fig. 6d, the coefficient estimate of *Displayed* for disgust in replies to political misleading posts is significantly positive ($coef. = 0.027$, $p < 0.01$; 95% CI: $[0.008, 0.046]$). However, the coefficient estimate of *Displayed* for disgust in replies to non-political misleading posts is statistically not significant ($coef. = 0.012$, $p = 0.473$; 95% CI: $[-0.020, 0.044]$). This indicates that the effect of community notes on increasing disgust in replies is only applicable to political misleading posts.
- As shown in Figs. 6e–6g, the coefficient estimates of *Displayed* for fear, joy, and sadness in replies to both political and non-political misleading posts are consistently not significant (each $p > 0.05$). This means that the display of community notes has no significant effects on fear, joy, and sadness in replies to either political or non-political misleading posts.
- In Fig. 6h, the coefficient estimates of *Displayed* for surprise in replies to both political ($coef. = -0.025$, $p < 0.05$; 95% CI: $[-0.050, -0.001]$) and non-political misleading posts ($coef. = -0.048$, $p < 0.001$; 95% CI: $[-0.074, -0.022]$) are significantly negative. This means that the display of community notes generally leads to a decrease of surprise in replies to both political and non-political misleading posts.

Additionally, we examine the associations between the sentiments and emotions in replies and the source sentiments and emotions across political and non-political misleading posts. Fig. 6a) shows that the coefficient estimate of *SourcePositive* is not statistically significant for positive sentiment in replies to political misleading posts ($coef. = -0.015$, $p = 0.350$; 95% CI: $[-0.047, 0.017]$), while it is significantly positive for positive sentiment in replies to non-political misleading posts ($coef. = 0.099$, $p < 0.001$; 95% CI: $[0.062, 0.136]$). Similarly, in Fig. 6f, the coefficient estimate of *SourceJoy* is statistically not significant for joy in replies to political misleading posts ($coef. = 0.008$, $p = 0.614$; 95% CI: $[-0.024, 0.041]$), while it is significantly positive for joy in replies to non-political misleading posts ($coef. = 0.062$, $p < 0.001$; 95% CI: $[0.031, 0.092]$). This suggests that the positive association between source positive sentiment (joy) in original posts and positive sentiment (joy) in replies disappears for political misleading posts. In contrast, for negative sentiment, anger, disgust, fear, sadness, and surprise in replies, there are consistently positive associations with their corresponding emotions in the source misleading posts, regardless of whether the posts are political or non-political (each $p < 0.05$, see Figs. 6b–6e and Figs. 6g–6h).

Finally, we find that the coefficient estimates of *HoursFromDisplay* for all sentiment/emotions in replies to either political or non-political misleading posts are consistently not significant (each

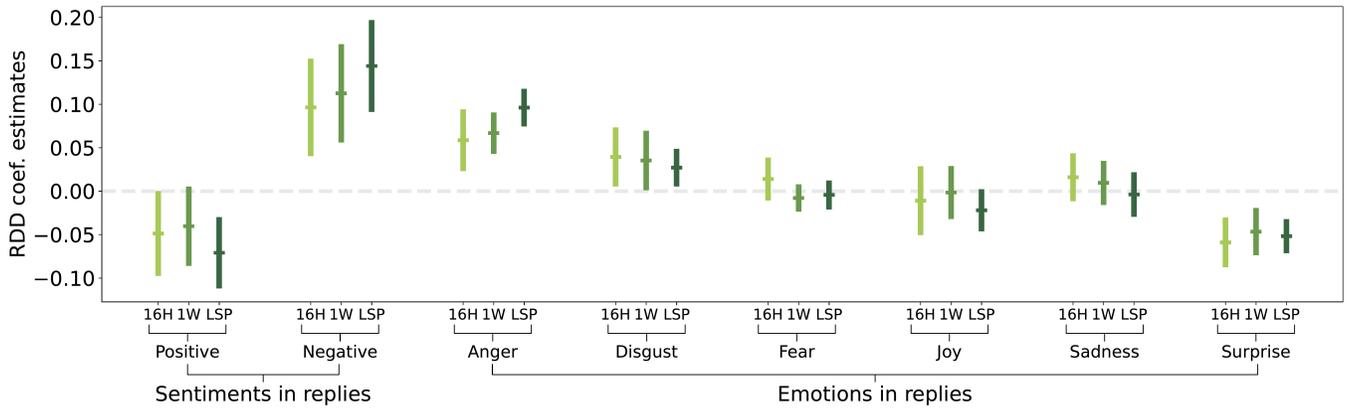


Figure 5: The RDD coefficient estimates (i. e., *Displayed*) for sentiments and emotions in replies across three different bandwidths. The three bandwidths are 16 hours (16H), 1 week (1W), and reply lifespan (LSP), respectively. Shown are mean values with error bars representing 95% CIs. Standard errors are clustered by source posts. The full estimation results are reported in Suppl. S4.

Table 2: SUEST tests for the equality of common coefficients across RDD models with different bandwidths. The tests are conducted separately for each sentiment/emotion in replies.

	(1)			(2)		
	16H vs. LSP			1W vs. LSP		
	χ^2	p	#Variables	χ^2	p	#Variables
<u>Sentiments in replies</u>						
Positive	8.47	0.132	5	10.74	0.057	5
Negative	3.97	0.554	5	4.30	0.507	5
<u>Emotions in replies</u>						
Anger	12.87	0.168	9	15.89	0.069	9
Disgust	9.03	0.434	9	8.25	0.509	9
Fear	6.23	0.717	9	8.16	0.518	9
Joy	8.10	0.524	9	8.45	0.490	9
Sadness	13.68	0.134	9	6.70	0.668	9
Surprise	6.49	0.690	9	7.54	0.581	9

$p > 0.05$; see Fig. 6). This suggests that the stability of sentiments and emotions over time remains robust, irrespective of whether the misleading posts are political or non-political.

Summary of findings: Taken together, we observe two aspects of our findings that are specific to political misleading posts: (i) The effects of community notes on the decrease of positive sentiment and the increase of disgust in replies only apply to political misleading posts; (ii) The transfer of positive sentiment, particularly joy, from source posts to replies is disrupted in discussions on political topics.

4.4 Analysis of Moral Outrage (RQ2.1 & RQ2.2)

Given that moral outrage is typically expressed through anger and disgust, we examine whether the display of community notes triggers the moral outrage of users in their replies to misleading posts. Analogous to previous research [28], we define the dependent variable moral outrage as $MoralOutrage = Anger \times Disgust$. We first examine the changes of moral outrage in replies to all misleading posts. As shown in Fig. 7a, the coefficient estimate of *Displayed* is significantly positive ($coef. = 0.067$, $p < 0.001$; 95%

CI: [0.050, 0.084]). This means that the display of community notes increases moral outrage in replies to misleading posts by 0.003 (95% CI: [0.002, 0.004]). Additionally, using the baseline of moral outrage during a period of 16 hours before note display (mean of 0.019, Fig. 7b), we find that the display of community notes triggers 16% more moral outrage in replies and increases it to a value of 0.023 (95% CI: [0.022, 0.023]).

Robustness with alternative classifiers: Previous research developed a dedicated classifier to predict moral outrage, namely, the Digital Outrage Classifier (DOC) [9]. As a robustness check, we repeat our analysis with this classifier and find that the coefficient estimate of *Displayed* remains significantly positive ($coef. = 0.067$, $p < 0.001$; 95% CI: [0.041, 0.092]). A SUEST test indicates no statistically significant difference in the effect of *Displayed* between the operationalization of moral outrage in our main analysis and the operationalization of moral outrage via the DOC classifier ($\chi^2 = 0.00$, $p = 0.957$). Additionally, anger and disgust are central components of other-condemning moral emotions [111]. Using a dedicated moral emotion classifier developed by prior research [67], we repeat our analysis and find that the effect of *Displayed*

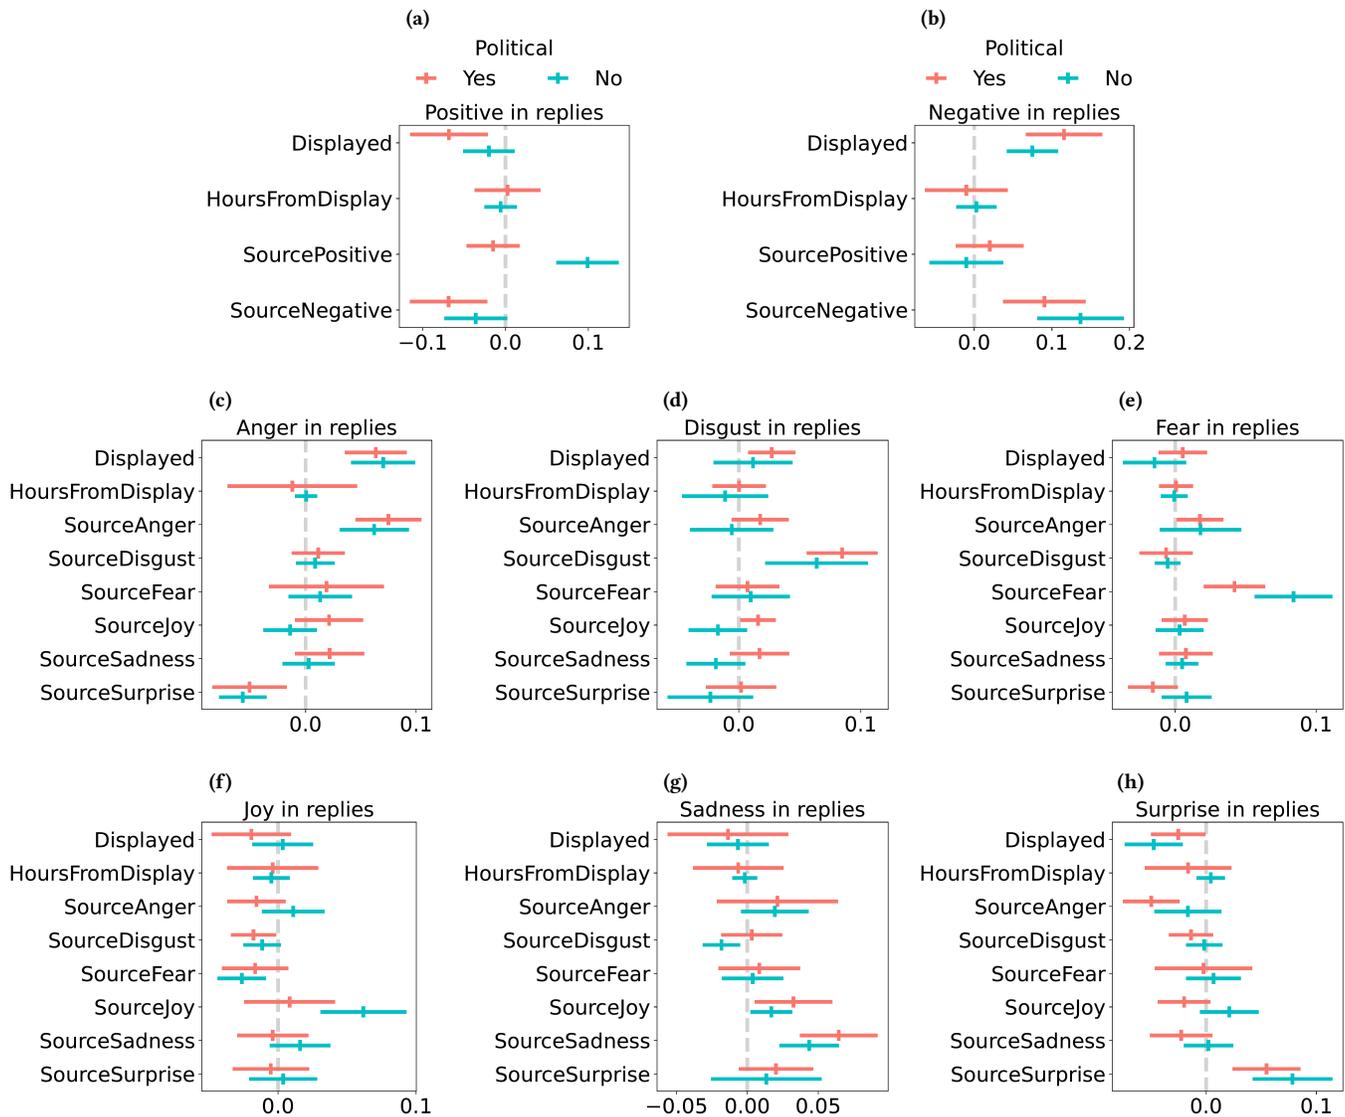


Figure 6: The estimated coefficients for the independent variables – *Displayed*, *HoursFromDisplay*, and source sentiments (or emotions) across political and non-political misleading posts. The independent variable *PostAge* is included during estimation but omitted in the visualization for better readability. Shown are mean values with error bars representing 95% CIs. Standard errors are clustered by source posts. The dependent variables are (a) positive sentiment in replies, (b) negative sentiment in replies, (c) anger in replies, (d) disgust in replies, (e) fear in replies, (f) joy in replies, (g) sadness in replies, and (h) surprise in replies, respectively. The full estimation results are reported in Suppl. S5.

on other-condemning moral emotions is also significantly positive ($coef. = 0.132, p < 0.001; 95\% CI: [0.103, 0.161]$). Altogether, these results consistently support that community notes trigger moral outrage beyond the isolated effects on anger and disgust (see details in Suppl. S6).

Analysis across political and non-political misleading posts: We further examine the effects of community notes display on moral outrage in replies to political and non-political misleading posts separately. As shown in Fig. 7c, the coefficient estimates of *Displayed* are significantly positive within both political ($coef. = 0.062, p <$

$0.001; 95\% CI: [0.037, 0.087]$) and non-political misleading posts ($coef. = 0.060, p < 0.001; 95\% CI: [0.036, 0.084]$). This indicates that the effects of community notes in triggering moral outrage are robust across political and non-political misleading posts. Additionally, we analyze the average levels of moral outrage in replies to political and non-political misleading posts before and after the display of community notes (Fig. 7d). Before note display, the averages of moral outrage in replies to political and non-political misleading posts are 0.025 and 0.020, respectively ($t = 63.709, p < 0.001$). After note display, these averages increase to 0.028 for political posts and

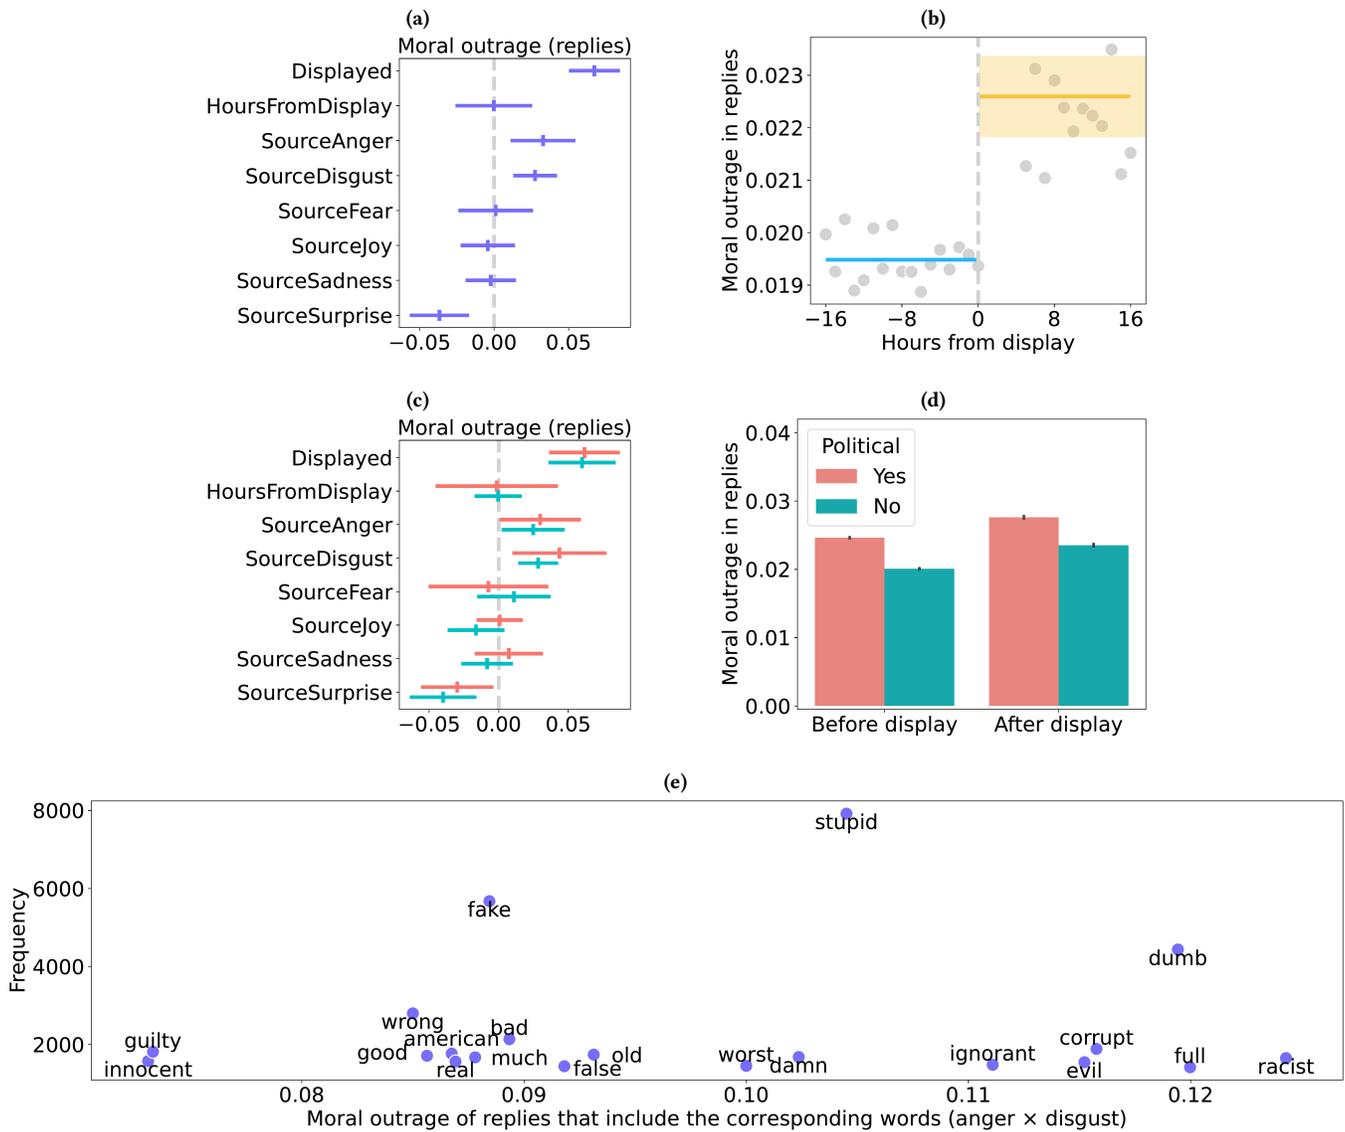


Figure 7: Moral outrage in replies to misleading posts increases after the display of community notes. (a) The estimated coefficients for the independent variables – *Displayed*, *HoursFromDisplay*, and source emotions. (b) The predicted effect of the display of community notes on moral outrage in replies to misleading posts. The error band represents 95% CIs. (c) The estimated coefficients for the independent variables – *Displayed*, *HoursFromDisplay* across political and non-political misleading posts. (d) The averages of moral outrage in replies to political and non-political misleading posts before and after the display of community notes. The error bars represent 95% CIs. (e) The twenty most frequent adjectives within replies with high moral outrage. For (a) and (c), the independent variable *PostAge* is included during estimation but omitted in the visualization for better readability. Shown are mean values with error bars representing 95% CIs. Standard errors are clustered by source posts. The full estimation results are reported in Suppl. S6.

0.024 for non-political posts ($t = 31.042$, $p < 0.001$). This suggests that the moral outrage in replies to political misleading posts is significantly higher than that in replies to non-political misleading posts, regardless of the display of community notes. However, the averages of moral outrage in the source posts between political (mean of 0.010) and non-political ones (mean of 0.009; $t = 0.575$,

$p = 0.566$) have no significant difference. This indicates that the heightened moral outrage in replies to political posts, compared to non-political posts, is not simply a transfer from the source posts' moral outrage.

Validation of targets of moral outrage: We conduct multiple tests to validate that the measured moral outrage in replies is

primarily directed at the misleading posts or their authors (rather than reflecting emotional reactions from users whose “side” has been labeled as misleading).

First, we consider two scenarios where replies may not be directed toward original posts: mentions of other users (not authors) in replies and mentions of community notes in replies. If the level of moral outrage in these two scenarios is lower than in replies directed solely toward original posts, this would suggest that the moral outrage in replies is primarily focused on misleading content. We find that the ratio of mentioning other users in replies after note display (10.4%) is lower than that in replies before note display (12.3%). Meanwhile, after note display, moral outrage in replies that do not mention others (mean of 0.028) is 180.0% more than that in replies that mention others (mean of 0.010; $t = -82.018$, $p < 0.001$). Additionally, after the display of community notes, replies may discuss community notes rather than the original misleading posts. Given this, we check how many replies mention community notes after their display through the keywords: “community note” and “CN.” We find that, out of 533,495 replies created after note display, only 2986 (0.6%) replies mention community notes. Moreover, moral outrage in replies that do not mention community notes (mean of 0.026) is 73.3% higher than the moral outrage in replies that mention community notes (mean of 0.015; $t = -11.791$, $p < 0.001$). Taken together, these validations suggest that the majority of replies, especially those with high moral outrage, are directed toward source misleading posts.

Second, we randomly select 100 misleading posts with displayed community notes and manually evaluate the reply with the highest outrage score after the display of community note for each selected post. Specifically, two authors of this study independently reviewed these 100 replies, arriving at identical results: all 100 replies are consistently directed toward the content of the misleading posts or their authors.

Third, we analyze the most frequently used words in replies that express moral outrage. We extract all words from replies posted after note display, focusing on those with a higher level of moral outrage (i. e., above the mean value). The most frequent word in these replies is “you,” often indicating a direct reference to the post author. Subsequently, we remove stop words and use Spacy, a Python library for advanced natural language processing, to extract the twenty most frequent adjectives. Then, we calculate the average level of moral outrage for each word based on the replies that include it. Fig. 7e shows that “stupid,” “fake,” and “dumb” are the most frequent adjectives in replies expressing moral outrage, with mean values of 0.105, 0.088, and 0.119, respectively. Additionally, words associated with higher levels of moral outrage (greater than 0.100) tend to be more critical (e. g., racist, dumb, and corrupt), while words associated with lower levels of moral outrage (below 0.100) are more related to the veracity of the content (e. g., wrong, fake, and false). This suggests that higher levels of moral outrage are associated with the use of more critical or aggressive words.

Summary of findings: Altogether, our analysis suggests that the display of community notes triggers moral outrage in replies to the corresponding misleading posts. Additionally, political misleading posts tend to receive more moral outrage compared to non-political misleading posts. However, the effects of community

notes on moral outrage in replies are robust across political and non-political misleading posts.

4.5 Robustness Checks

We conduct a wide range of additional robustness checks to ensure the reliability of our findings:

- **Multicollinearity check:** We check for possible multicollinearity issues using Variance Inflation Factors (VIFs). The VIFs are all close to one; and, thus well below the critical threshold of four (see details in Suppl. S7.1).
- **Analysis with lagged-dependent variables:** To address the possibility that users replying to original posts may be influenced by sentiments and emotions in previous direct replies, we conduct an auto-correlation test and repeat our analysis with lagged-dependent variables. The results remain robust, consistently supporting our main findings (see details in Suppl. S7.2).
- **Shorter bandwidth and placebo test:** To further ensure that the increase in moral outrage in replies to misleading posts is specific to the display of community notes, we repeat our analysis within a shorter bandwidth, i. e., 2 hours, around the display of community notes and conduct a placebo test at a different cut-off point before note display for comparison. We find that all results are robust and consistent with our main analysis (see details in Suppl. S7.3).
- **Analysis without reply restriction:** We repeat the analysis without the restriction of 5 replies before and after note display. The results remain robust and consistent with our main findings (see details in Suppl. S7.4).
- **Month-year fixed effects:** We incorporate month-year fixed effects into our regression models and repeat our analysis. The results remain robust and consistent with our main findings (see details in Suppl. S7.5).
- **Validation with dictionary-based approaches:** We repeat our analysis using sentiment and emotion lexicons. Specifically, we use the VADER lexicon for sentiments and the NRC lexicon for emotions [60, 81]. The results remain robust across machine learning models and lexicons (see details in Suppl. S7.6).
- **Role of sentiments and emotions in community notes:** To explore whether sentiments and emotions in community notes can potentially moderate the effects of community notes display, we incorporate sentiments and emotions in community notes and their interactions with *Displayed* into our regression models. The coefficient estimates of *Displayed* across the dependent variables remain robust. Additionally, the majority of the coefficient estimates for the interactions between sentiments (emotions) in community notes and *Displayed* are not statistically significant. This suggests that users’ emotional reactions are not significantly affected by the sentiments or emotions in community notes (see details in Suppl. S7.7).
- **Robustness across helpfulness scores:** Misleading posts that garner significant public concern may be more likely to be quoted by influential accounts on X, thereby reaching a broader audience. This increased visibility may introduce confounding factors related to the helpfulness scores of community notes. For example, counter-partisan accounts can encourage other users to rate a community note as helpful, thereby increasing the helpfulness

score of the note [3]. To account for such potential confounding factors, we incorporate *NoteScore* and its interaction with *Displayed* into our regression models. We find that the effect of community notes is robust across different helpfulness scores (see details in Suppl. S7.8).

- **Effect of community notes over the program’s development:** Given that the Community Notes program continuously evolves, we incorporate *MFRO* (i. e., months from the roll-out of Community Notes program) and its interaction with *Displayed* into regression models to investigate how the effect of community notes changes over time from the roll-out of the program. We find that the effect of community notes on our dependent variables remains stable (see details in Suppl. S7.9).

4.6 Summary of Main Findings

In summary, we conduct a comprehensive *causal* analysis on the role of community notes in triggering sentiments, emotions, and moral outrage in replies to fact-checked social media posts. Furthermore, we extend our analysis by conducting sensitivity analysis across political and non-political misleading posts. Our main findings are as follows:

- The display of community notes triggers 7.3% more negativity, 13.2% more anger, and 4.7% more disgust in replies to misleading posts (RQ1.1).
- The effects of community notes on negativity and anger are robust across political and non-political misleading posts, while the effect on disgust only applies to political misleading posts (RQ1.2).
- The display of community notes triggers 16% more moral outrage (RQ2.1), and this is robust across political and non-political misleading posts (RQ2.2).

In addition to our main findings, we find the evidence of emotion transfers and time-independence for sentiments and emotions in online social networks:

- The sentiments and emotions in replies to misleading posts are positively correlated with the sentiments and emotions of the corresponding source posts (with very few exceptions).
- The sentiments and emotions in replies are stable over time if the posts are not subject to fact-checking (i. e., do not receive a displayed community note).

5 Discussion

5.1 Relevance

The spread of online misinformation has become a significant challenge that social media platforms have to tackle. However, traditional (expert-based) fact-checking approaches have limitations regarding their scalability and face trust issues among the user base. As a remedy, crowdsourced fact-checking represents a promising approach to identify misinformation at scale and increase users’ trust in fact-checks. As the first large-scale attempt of crowdsourced fact-checking approach on a major social media platform, the “Community Notes” program has been shown to be effective in producing trustworthy community fact-checks and reducing engagement with misleading posts on X [19, 36]. It provides a reference for other social media platforms to implement similar crowdsourced fact-checking approaches (e. g., YouTube [128]). However, to improve

and extend the “Community Notes” program safely, it is important to understand how community notes affect the discussion environment after they are displayed on misleading posts, i. e., how users’ emotional expressions change after they are informed about the falsehood. To this end, our study causally examines the changes in emotional expressions in replies to misleading posts before and after the display of community notes. As detailed in the next sections, our findings provide important insights into the effects of (community-based) fact-checking on the discussion environment on social media and may help to improve the design of crowdsourced fact-checking platforms.

5.2 Interpretations

5.2.1 Community notes trigger negativity, anger, and disgust in replies to misleading posts. We find that the display of community notes leads users to post more negative replies, which is primarily due to heightened expressions of anger and disgust. These observations align with previous survey and lab studies in the field of psychology that people often report anger and disgust in response to moral violations of community, autonomy, and divinity ethics [82, 119]. In the context of our study, the display of community notes signals to users that the posts they are exposed to spread misinformation, which may thereby trigger their perception of a moral violation. Anger and disgust, though both negative, differ in their aggressive tendencies when responding to moral violations. Anger is typically associated with high-cost, direct aggression (e. g., physical violence or confrontational responses), whereas disgust is typically associated with less costly indirect aggression (e. g., gossip and social exclusion). Additionally, previous research suggests that people tend to express more anger than disgust when moral violations directly affect themselves, while they express more disgust than anger when moral violations target others [82]. Quantitatively, our findings reveal that the display of community notes increases anger in replies to misleading posts by 13.2%, compared to a 4.7% increase in disgust. Thus, users are more likely to express anger than disgust, which may indicate that users perceive the misleading content to primarily impact their own interests (rather than others’). Notably, our sensitivity analysis further reveals that the effect of community notes in increasing disgust is statistically significant only in replies to political misleading posts. This could partially explain why the display of community notes is less effective in reducing the spread of political misleading posts, compared to non-political ones, which has been observed in a previous study [19, 20].

5.2.2 Community notes trigger moral outrage in replies to misleading posts. When anger and disgust are mixed, they generate “moral outrage,” a more intense emotional reaction to perceived moral transgressions than either anger or disgust alone [9, 28, 104]. Our findings show that the display of community notes increases moral outrage by 16% in replies to misleading posts, a greater increase than those observed for anger (13.2%) and disgust (4.7%) individually. This empirical evidence underscores the significant role of moral outrage in users’ responses to the moral transgression of spreading misleading posts. Specifically, our findings suggest that displaying community notes on misleading posts is a strong stimulus that calls attention to the moral norm violation, thereby motivating

expressions of moral outrage. Notably, research suggests that if users experience moral violations personally without such external stimuli, their intentions to express moral outrage might decrease [28]. This phenomenon might be explained by social conformity that users are inclined to conform to the majority's opinions on the quality of online posts [27, 122]. The display of community notes signals consensus among diverse users, thereby potentially encouraging exposed users' conformity and subsequent correcting actions.

The expression of moral outrage to spreading misinformation is a double-edged sword that can both mitigate the harm of online misinformation and exacerbate the polarization of online communities. On the one hand, moral outrage can strengthen debunking effects on exposed users and pressure authors to correct or delete their misleading posts [19, 88]. Previous research has shown that the display of community notes significantly increases the likelihood of deleting misleading posts by their authors [19]. Moreover, those who attempt to spread misinformation can expect strong negative reactions from the user base once their content is debunked, which may help to discipline users and encourage more responsible behavior on social media. On the other hand, social media platforms also need to maintain a healthy, inclusive, and balanced online environment. However, online moral outrage can deepen social divides and fuel political polarization [9, 28, 37]. Additionally, the corrections and moral judgments from peer users may backfire and make authors of misleading posts feel defensive, promoting their subsequent shares of low-quality, partisan, and toxic content [83]. In this light, while moral outrage can be a powerful tool for combating misinformation, it also risks escalating conflicts and entrenching divisions within online communities. This risk may partially explain why community notes lead to smaller reductions in replies compared to reposts [97]. The potential negative effects of moral outrage triggered by community notes require a further rigorous examination to find a trade-off between mitigating the spread of misinformation and maintaining a positive discussion atmosphere.

5.2.3 Emotion transfers and stable emotional responses over time. Our study provides strong empirical evidence supporting the idea that connected users on social media platforms express similar sentiments/emotions. Specifically, we find that all sentiments and emotions in misleading posts, especially non-political ones, are effectively transferred to subsequent sentiments and emotions in replies. This emotion transfer phenomenon can be explained primarily through two mechanisms: emotion contagion and emotion alignment (i. e., homophily) among connected users. Here, emotion contagion refers to the spreading effect of emotions from one user to another [51, 100], and emotion homophily indicates that like-minded users tend to flock together, potentially forming echo chambers that dominate online interactions and reinforce biases and attitude polarization [25]. Our findings further reveal that the intensity of sentiments and emotions in replies remains stable over time, showing no signs of fading. Moreover, misleading posts fact-checked via community notes are predominantly characterized by negativity. The persistence of emotion transfers may further exacerbate these dynamics [22, 30, 133].

5.3 Implications for Further Design of Community-Based Fact-Checking

Community-based fact-checking is a promising approach and can effectively reduce engagement with misinformation on social media [2, 19, 36, 62, 131]. However, our study finds that communicating falsehoods to users via community fact-checks can also trigger moral outrage in response to the perceived moral violation of spread misinformation. This heightened moral outrage can have dual impacts. On the one hand, it may strengthen the effectiveness of community-based fact-checking by pressuring (disciplining) users to reconsider their actions. On the other hand, however, it may also backfire by increasing tension and polarization on social media [83].

The dual role of moral outrage highlights the need to minimize its negative consequences without undermining its positive effects in the design and implementation of community-based fact-checking to maintain health and constructive dialogue on social media. Previous research suggests that temporarily deactivating social media accounts can reduce polarization on policy issues [95]. Additionally, a “cooldown” period for highly active and polarized threads has been proposed to ensure a deliberately interactive environment [39]. Given this, social media platforms might consider implementing a *temporary freezing period* for misleading posts once they receive helpful community fact-checks. During this period – e. g., one hour – the misleading posts would be invisible to others, while the authors are notified, giving them time to delete or correct the misleading content. If they do not respond within the given timeframe, the community fact-checks with associated misleading posts would then be publicly displayed on the platform. Similarly, platforms may consider *temporarily limiting the ability to comment* on posts flagged by community fact-checks. This could take the form of a complete ban on posting responses during this period or a personal delay, where written comments are visible to others only after a certain period. In general, platforms may need to emphasize *community guidelines* in heated/morally charged discussions (e. g., labels/brief prompts that appear when a user is about to post a highly emotional response; or additional monitoring by moderation teams) [98].

Importantly, however, any such design changes must be carefully evaluated through further research to assess their potential benefits and drawbacks. A core motivation behind community-based fact-checking is to reduce censorship on social media platforms and empower users with greater autonomy to improve trust in the fact-checking process [36]. Therefore, platforms must strike a delicate balance between fostering constructive discourse and preserving user autonomy to sustain trust and engagement in community-based fact-checking.

5.4 Limitations and Future Research

Our study has several limitations that could open up potential avenues for future research. First, our study is limited to a 4-month period since the roll-out of the “Community Notes” program. As the program is continually evolving, it would be valuable to examine the effects of community notes on sentiments, emotions, and moral

outrage over a longer observation timeframe.¹³ Furthermore, while we conduct extensive robustness checks to ensure the reliability of our findings, it is challenging for a quasi-experimental study to fully isolate the effect of community notes from other potential concurrent misinformation interventions or external factors. Therefore, future research is encouraged to validate our findings within controlled experimental designs.

Second, on the date of data collection, the “Community Notes” program was primarily well-established in the U.S., and the fact-checked posts were mainly in English. Consequently, we only consider English posts in our study. In the future, a multilingual and cross-cultural evaluation would be beneficial to get a more comprehensive understanding of how community notes impact users’ emotional reactions across languages and cultures. Additionally, our sensitivity analysis is limited to general political and non-political misleading posts. Future research could expand this scope by examining more fine-grained topics (e. g., health, business & economics, celebrities, etc.).

Third, our measurement of moral outrage cannot definitively ascertain whether it is directed toward misleading posts or their authors. To validate our approach, we analyze two special cases: mentions of other users and mentions of community notes in replies. We find that the moral outrage in replies that mention other users or community notes is significantly lower than the moral outrage in replies that solely react to the misleading posts and do not mention other users or community notes. Additionally, replies with high moral outrage frequently use aggressive words (e. g., “stupid,” “dumb”). All of these analyses support the notion that community notes trigger moral outrage and critical remarks to misleading posts or their authors. Future research could further validate our findings through more advanced subject-aware emotion detection models.

Finally, it has been found that the display of community notes increases the probability of post deletion by authors [20]. A possible contributing factor may be that expressions of moral outrage pressure authors to delete their misleading posts. This notion could be validated through the examination of moral outrage in replies to deleted and non-deleted posts. However, deleted misleading posts and corresponding replies cannot be collected through X’s API. As a remedy, future research could conduct lab experiments to causally examine whether expressions of moral outrage increase authors’ intentions to delete their misleading posts. Additionally, it would be interesting to further examine how the authors of misleading posts adapt their future behaviors on social media.

6 Conclusion

Community-based fact-checking is a promising approach to reduce engagement with misinformation on social media. Yet, an understanding of how community fact-checks affect users’ emotional perception of misleading posts and their authors was missing. In this study, we analyze a large-scale panel dataset of replies to misleading posts that have been fact-checked via community notes. Our findings provide strong causal evidence that community fact-checks trigger moral outrage in replies to misleading posts. This

suggests that social media users perceive spreading misinformation as moral transgressions, and that spreaders of misinformation should expect negative responses once their content is debunked. Our study highlights the need for further examination of both the positive and negative consequences of such emotional backlash.

7 Ethics Statement

This research has received ethical approval from the Ethics Review Panel of the University of Luxembourg (ref. ERP 23-053 REMEDIS). All analyses are based on publicly available data. To respect privacy, we explicitly do not publish usernames in our paper and only report aggregate results. We declare no competing interests.

Acknowledgments

This research is supported by the Luxembourg National Research Fund (FNR) and Belgian National Fund for Scientific Research (FNRS), as part of the project REgulatory Solutions to MitigatE DIS-information (REMEDIS), ref. INTER_FNRS_21_16554939_REMEDIS. Furthermore, this research is supported by a research grant from the German Research Foundation (DFG grant 492310022).

References

- [1] Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of Economic Perspectives* 31, 2 (2017), 211–236.
- [2] Jennifer Allen, Antonio A Arechar, Gordon Pennycook, and David G Rand. 2021. Scaling up fact-checking using the wisdom of crowds. *Science Advances* 7, 36 (2021), eabf4393.
- [3] Jennifer Allen, Cameron Martel, and David G Rand. 2022. Birds of a feather don’t fact-check each other: Partisanship and the evaluation of news in Twitter’s Birdwatch crowdsourced fact-checking program. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [4] Sacha Altay, Manon Berriche, and Alberto Acerbi. 2023. Misinformation on misinformation: Conceptual and methodological challenges. *Social Media + Society* 9, 1 (2023), 20563051221150412.
- [5] Joshua D Angrist and Jörn-Steffen Pischke. 2009. *Mostly harmless econometrics: An empiricist’s companion*. Princeton University Press.
- [6] Sinan Aral and Dean Eckles. 2019. Protecting elections from social media manipulation. *Science* 365, 6456 (2019), 858–861.
- [7] Jonah Berger and Katherine L Milkman. 2012. What makes online content viral? *Journal of Marketing Research* 49, 2 (2012), 192–205.
- [8] Md Momen Bhuiyan, Amy X Zhang, Connie Moon Sehat, and Tanushree Mitra. 2020. Investigating differences in crowdsourced news credibility assessment: Raters, tasks, and expert criteria. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–26.
- [9] William J Brady, Killian McLoughlin, Tuan N Doan, and Molly J Crockett. 2021. How social learning amplifies moral outrage expression in online social networks. *Science Advances* 7, 33 (2021), eabe5641.
- [10] William J Brady, Killian L McLoughlin, Mark P Torres, Kara F Luo, Maria Gendron, and MJ Crockett. 2023. Overperception of moral outrage in online social networks inflates beliefs about intergroup hostility. *Nature Human Behaviour* 7, 6 (2023), 917–927.
- [11] William J Brady, Julian A Wills, John T Jost, Joshua A Tucker, and Jay J Van Bavel. 2017. Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences* 114, 28 (2017), 7313–7318.
- [12] William Brannon and Deb Roy. 2024. The speed of news in Twitter (X) versus radio. *Scientific Reports* 14, 1 (2024), 11939.
- [13] Ceren Budak, Brendan Nyhan, David M Rothschild, Emily Thorson, and Duncan J Watts. 2024. Misunderstanding the harms of online misinformation. *Nature* 630, 8015 (2024), 45–53.
- [14] Sabur Butt, Shakshi Sharma, Rajesh Sharma, Grigori Sidorov, and Alexander Gelbukh. 2022. What goes on inside rumour and non-rumour tweets and their reactions: A psycholinguistic analyses. *Computers in Human Behavior* 135 (2022), 107345.
- [15] Dominik Bär, Nicolas Pröllochs, and Stefan Feuerriegel. 2023. New threats to society from free-speech social media platforms. *Commun. ACM* 66, 10 (2023), 37–40.
- [16] Jose Camacho-collados, Kiamehr Rezaee, Talayah Riahi, Asahi Ushio, Daniel Loureiro, Dimosthenis Antypas, Joanne Boisson, Luis Espinosa Anke, Fangyu Liu, and Eugenio Martínez Cámara. 2022. TweetNLP: Cutting-edge natural

¹³Notably, we collected our dataset via X’s Academic Research API before it was deprecated. Currently, it is costly to collect such a high volume of data from X.

- language processing for social media. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. 38–49.
- [17] Matias D Cattaneo and Rocio Titiunik. 2022. Regression discontinuity designs. *Annual Review of Economics* 14, 1 (2022), 821–851.
- [18] Yuwei Chuai, Yutian Chang, and Jichang Zhao. 2022. What really drives the spread of COVID-19 Tweets: A revisit from perspective of content. In *2022 IEEE 9th International Conference on Data Science and Advanced Analytics*. 1–10.
- [19] Yuwei Chuai, Moritz Pilarski, Gabriele Lenzini, and Nicolas Pröllochs. 2024. Community notes reduce the spread of misleading posts on X. *OSF* (2024).
- [20] Yuwei Chuai, Moritz Pilarski, Thomas Renault, David Restrepo-Amariles, Aurore Troussel-Clément, Gabriele Lenzini, and Nicolas Pröllochs. 2024. Community-based fact-checking reduces the spread of misleading posts on social media. *ArXiv* (2024).
- [21] Yuwei Chuai, Haoye Tian, Nicolas Pröllochs, and Gabriele Lenzini. 2024. Did the roll-out of Community Notes reduce engagement with misinformation on X/Twitter? *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW2 (2024), 1–52.
- [22] Yuwei Chuai and Jichang Zhao. 2022. Anger can make fake news viral online. *Frontiers in Physics* 10 (2022), 970174.
- [23] Yuwei Chuai, Jichang Zhao, and Gabriele Lenzini. 2024. From news sharers to post viewers: How topic diversity and conspiracy theories shape engagement with misinformation during a health crisis? *ArXiv* (2024).
- [24] Yuwei Chuai, Jichang Zhao, Nicolas Pröllochs, and Gabriele Lenzini. 2025. Is fact-checking politically neutral? Asymmetries in how U.S. fact-checking organizations pick up false statements mentioning political elites. In *Proceedings of the International AAI Conference on Web and Social Media*, Vol. forthcoming.
- [25] Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2021. The echo chamber effect on social media. *Proceedings of the National Academy of Sciences* 118, 9 (2021), e2023301118.
- [26] Katherine Clayton, Spencer Blair, Jonathan A Busam, Samuel Forstner, John Glance, Guy Green, Anna Kawata, Akhila Kovvuri, Jonathan Martin, Evan Morgan, et al. 2020. Real solutions for fake news? Measuring the effectiveness of general warnings and fact-check tags in reducing belief in false stories on social media. *Political Behavior* 42 (2020), 1073–1095.
- [27] Jonas Colliander. 2019. “This is fake news”: Investigating the role of conformity to other users’ views when commenting on and spreading disinformation in social media. *Computers in Human Behavior* 97 (2019), 202–215.
- [28] Molly J Crockett. 2017. Moral outrage in the digital age. *Nature Human Behaviour* 1, 11 (2017), 769–771.
- [29] Ross Dahlke and Jeffrey Hancock. 2024. Effects of misinformation on election beliefs: Disentangling motivated reasoning from selective exposure. *OSF* (2024).
- [30] Michela Del Vicario, Gianna Vivaldo, Alessandro Bessi, Fabiana Zollo, Antonio Scala, Guido Caldarelli, and Walter Quattrociocchi. 2016. Echo chambers: Emotional contagion and group polarization on facebook. *Scientific Reports* 6, 1 (2016), 37825.
- [31] Deloitte. 2021. Are younger generations moving away from traditional news sources? <https://www.forbes.com/sites/kenrapoza/2017/02/26/can-fake-news-impact-the-stock-market/>.
- [32] Joan Donovan. 2020. Social-media companies must flatten the curve of misinformation. *Nature* (2020).
- [33] Chiara Patricia Drolsbach and Nicolas Pröllochs. 2023. Diffusion of community fact-checked misinformation on Twitter. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (2023), 1–22.
- [34] Chiara Patricia Drolsbach and Nicolas Pröllochs. 2024. Content moderation on social media in the EU: Insights from the DSA Transparency Database. In *Companion Proceedings of the ACM on Web Conference 2024*. 939–942.
- [35] Chiara Patricia Drolsbach and Nicolas Pröllochs. 2023. Believability and harmfulness shape the virality of misleading social media posts. In *Proceedings of the ACM Web Conference 2023*. 4172–4177.
- [36] Chiara Patricia Drolsbach, Kirill Solovev, and Nicolas Pröllochs. 2024. Community notes increase trust in fact-checking on social media. *PNAS Nexus* 3, 7 (2024), pgae217.
- [37] Ullrich Ecker, Jon Roozbeek, Sander van der Linden, Li Qian Tay, John Cook, Naomi Oreskes, and Stephan Lewandowsky. 2024. Misinformation poses a bigger threat to democracy than you might think. *Nature* 630, 8015 (2024), 29–32.
- [38] Ullrich KH Ecker, Stephan Lewandowsky, John Cook, Philipp Schmid, Lisa K Fazio, Nadia Brashier, Panayiota Kendeou, Emily K Vraga, and Michelle A Amazeen. 2022. The psychological drivers of misinformation belief and its resistance to correction. *Nature Reviews Psychology* 1, 1 (2022), 13–29.
- [39] Alexandros Efstratiou and Emiliano De Cristofaro. 2022. Adherence to misinformation on social media through socio-cognitive and group-based processes. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–35.
- [40] Paul Ekman. 1992. An argument for basic emotions. *Cognition & Emotion* 6, 3–4 (1992), 169–200.
- [41] Paul Ekman, Wallace V Friesen, and Phoebe Ellsworth. 2013. *Emotion in the human face: Guidelines for research and an integration of findings*. Vol. 11. Elsevier.
- [42] Adam M. Enders, Joseph Uscinski, Casey Klofstad, and Justin Stoler. 2022. On the relationship between conspiracy theory beliefs, misinformation, and vaccine hesitancy. *PLoS ONE* 17, 10 (2022), e0276082.
- [43] Ziv Epstein, Gordon Pennycook, and David Rand. 2020. Will the crowd game the algorithm? Using layperson judgments to combat misinformation on social media by downranking distrusted sources. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [44] Morten W Fagerland. 2012. T-tests, non-parametric tests, and large studies—a paradox of statistical practice? *BMC Medical Research Methodology* 12 (2012), 1–7.
- [45] Stefan Feuerriegel, Renée DiResta, Josh A Goldstein, Srijan Kumar, Philipp Lorenz-Spreen, Michael Tomz, and Nicolas Pröllochs. 2023. Research can help to tackle AI-generated disinformation. *Nature Human Behaviour* 7, 11 (2023), 1818–1821.
- [46] Stefan Feuerriegel, Abdurahman Maarouf, Dominik Bär, Dominique Geissler, Jonas Schweisthal, Nicolas Pröllochs, Claire E Robertson, Steve Rathje, Jochen Hartmann, Saif M Mohammad, Oded Netzer, Alexandra Siegel, Barbara Plank, and Van Bavel Jay. 2025. Using natural language processing to analyse text data in behavioural science. *Nature Reviews Psychology* forthcoming (2025).
- [47] Daniela Flamini. 2019. Most Republicans don’t trust fact-checkers, and most Americans don’t trust the media. <https://www.poynter.org/ifcn/2019/most-republicans-dont-trust-fact-checkers-and-most-americans-dont-trust-the-media/>.
- [48] World Economic Forum. 2024. Global risks report. <https://www.weforum.org/publications/global-risks-report-2024/>.
- [49] Felix Gaisbauer, Armin Pournaki, Sven Banisch, and Eckehard Olbrich. 2021. Ideological differences in engagement in public debate on Twitter. *PLoS ONE* 16, 3 (2021), e0249241.
- [50] Dilrukshi Gamage, Piyush Ghasiya, Vamshi Bonagiri, Mark E Whiting, and Kazutoshi Sasahara. 2022. Are deepfakes concerning? Analyzing conversations of deepfakes on reddit and exploring societal implications. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [51] Amit Goldenberg and James J Gross. 2020. Digital emotion contagion. *Trends in Cognitive Sciences* 24, 4 (2020), 316–328.
- [52] Jon Green, William Hobbs, Stefan McCabe, and David Lazer. 2022. Online engagement with 2020 election misinformation and turnout in the 2021 Georgia runoff election. *Proceedings of the National Academy of Sciences* 119, 34 (2022), e2115900119.
- [53] Jonathan Haidt et al. 2003. The moral emotions. *Handbook of Affective Sciences* 11, 2003 (2003), 852–870.
- [54] Md Mahfuzul Haque, Mohammad Yousuf, Ahmed Shatil Alam, Pratyasha Saha, Syed Ishtiaque Ahmed, and Naemul Hassan. 2020. Combating misinformation in Bangladesh: Roles and responsibilities as perceived by journalists, fact-checkers, and users. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–32.
- [55] P Sol Hart, Sedona Chinn, and Stuart Soroka. 2020. Politicization and polarization in COVID-19 news coverage. *Science Communication* 42, 5 (2020), 679–697.
- [56] Jochen Hartmann. 2022. Emotion english DistilRoBERTa-base. <https://huggingface.co/j-hartmann/emotion-english-distilroberta-base/>.
- [57] Amelia Hassoun, Ian Beacock, Sunny Consolvo, Beth Goldberg, Patrick Gage Kelley, and Daniel M Russell. 2023. Practicing information sensibility: How Gen Z engages with online information. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [58] Catherine Hausman and David S Rapson. 2018. Regression discontinuity in time: Considerations for empirical applications. *Annual Review of Resource Economics* 10, 1 (2018), 533–552.
- [59] Hendrik Heuer and Elena Leah Glassman. 2022. A comparative evaluation of interventions against misinformation: Augmenting the WHO checklist. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–21.
- [60] Clayton Hutto and Eric Gilbert. 2014. VADER: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the International AAI Conference on Web and Social Media*, Vol. 8. 216–225.
- [61] Guido W Imbens and Thomas Lemieux. 2008. Regression discontinuity designs: A guide to practice. *Journal of Econometrics* 142, 2 (2008), 615–635.
- [62] Farnaz Jahanbakhsh and David R Karger. 2024. A browser extension for in-place signaling and assessment of misinformation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–21.
- [63] Johannes Jakubik, Michael Vössing, Nicolas Pröllochs, Dominik Bär, and Stefan Feuerriegel. 2023. Online emotions during the storming of the US Capitol: Evidence from the social media network Parler. In *Proceedings of the International AAI Conference on Web and Social Media*, Vol. 17. 423–434.
- [64] Prerna Juneja, Md Momen Bhuiyan, and Tanushree Mitra. 2023. Assessing enactment of content regulation policies: A post hoc crowd-sourced audit of election misinformation on YouTube. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–22.

- [65] Naveena Karusala and Richard Anderson. 2022. Towards conviviality in navigating health information on social media. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [66] Luke Keele and Nathan J Kelly. 2006. Dynamic models for dynamic theories: The ins and outs of lagged dependent variables. *Political Analysis* 14, 2 (2006), 186–205.
- [67] Jaehong Kim, Chaeyoon Jeong, Seongchan Park, Meeyoung Cha, and Wonjae Lee. 2024. How do moral emotions shape political participation? A cross-cultural analysis of online petitions using language models. In *Findings of the Association for Computational Linguistics ACL 2024*. 16274–16289.
- [68] Anastasia Kozyreva, Stefan M Herzog, Stephan Lewandowsky, Ralph Hertwig, Philipp Lorenz-Spreen, Mark Leiser, and Jason Reifer. 2023. Resolving content moderation dilemmas between free speech and harmful misinformation. *Proceedings of the National Academy of Sciences* 120, 7 (2023), e2210666120.
- [69] Adam DI Kramer, Jamie E Guillory, and Jeffrey T Hancock. 2014. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* 111, 24 (2014), 8788–8790.
- [70] David MJ Lazer, Matthew A Baum, Yochai Benkler, Adam J Berinsky, Kelly M Greenhill, Filippo Menczer, Miriam J Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, and others. 2018. The science of fake news. *Science* 359, 6380 (2018), 1094–1096.
- [71] Sian Lee, Aiping Xiong, Haeseung Seo, and Dongwon Lee. 2023. “Fact-checking” fact checkers: A data-driven approach. *Harvard Kennedy School Misinformation Review* (2023).
- [72] Jennifer S Lerner, Ye Li, Piercarlo Valdesolo, and Karim S Kassam. 2015. Emotion and decision making. *Annual Review of Psychology* 66, 1 (2015), 799–823.
- [73] Daniel Loureiro, Francesco Barbieri, Leonardo Neves, Luis Espinosa Anke, and Jose Camacho-collados. 2022. TimeLMs: Diachronic language models from Twitter. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. 251–260.
- [74] Thomas Lumley, Paula Diehr, Scott Emerson, and Lu Chen. 2002. The importance of the normality assumption in large public health data sets. *Annual Review of Public Health* 23, 1 (2002), 151–169.
- [75] Cameron Martel, Gordon Pennycook, and David G Rand. 2020. Reliance on emotion promotes belief in fake news. *Cognitive Research: Principles and Implications* 5 (2020), 1–20.
- [76] Melinda McClure Haughey, Martina Povolo, and Kate Starbird. 2022. Bridging contextual and methodological gaps on the “misinformation beat”: Insights from journalist-researcher collaborations at speed. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [77] Killian L McLoughlin, William J Brady, Aden Goolsbee, Ben Kaiser, Kate Klonick, and MJ Crockett. 2024. Misinformation exploits outrage to spread online. *Science* 386, 6725 (2024), 991–996.
- [78] Meta. 2024. Meta’s third-party fact-checking program. <https://www.facebook.com/formedia/mjp/programs/third-party-fact-checking>
- [79] Nicholas Micallef, Vivienne Armacost, Nasir Memon, and Sameer Patil. 2022. True or false: Studying the work practices of professional fact-checkers. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW1 (2022), 1–44.
- [80] Nicholas Micallef, Bing He, Srijan Kumar, Mustaque Ahamad, and Nasir Memon. 2020. The role of the crowd in countering misinformation: A case study of the COVID-19 infodemic. In *2020 IEEE International Conference on Big Data*. 748–757.
- [81] Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a word–emotion association lexicon. *Computational Intelligence* 29, 3 (2013), 436–465.
- [82] Catherine Molho, Joshua M Tybur, Ezgi Güler, Daniel Balliet, and Wilhelm Hofmann. 2017. Disgust and anger relate to different aggressive responses to moral violations. *Psychological Science* 28, 5 (2017), 609–619.
- [83] Mohsen Mosleh, Cameron Martel, Dean Eckles, and David Rand. 2021. Perverse downstream consequences of debunking: Being corrected by another user for posting false political news increases subsequent sharing of low quality, partisan, and toxic content in a Twitter field experiment. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [84] Dominik Neumann and Nancy Rhodes. 2024. Morality in social media: A scoping review. *New Media & Society* 26, 2 (2024), 1096–1126.
- [85] Society of Professional Journalists. 2014. Code of ethics. <https://www.spj.org/ethicscode.asp>
- [86] Mathias Osmundsen, Alexander Bor, Peter Bjerregaard Vahlstrup, Anja Bechmann, and Michael Bang Petersen. 2021. Partisan polarization is the primary psychological motivation behind political fake news sharing on Twitter. *American Political Science Review* 115, 3 (2021), 999–1015.
- [87] Bo Pang, Lillian Lee, et al. 2008. Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval* 2, 1–2 (2008), 1–135.
- [88] Kun Peng, Yu Zheng, Yuewei Qiu, and Qingrui Li. 2023. Rage of righteousness: Anger’s role in promoting debunking effects during the COVID-19 pandemic. *Computers in Human Behavior* 148 (2023), 107896.
- [89] Gordon Pennycook and David G Rand. 2019. Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences* 116, 7 (2019), 2521–2526.
- [90] Francesco Pierri, Brea L. Perry, Matthew R. DeVerna, Kai-Cheng Yang, Alessandro Flammini, Filippo Menczer, and John Bryden. 2022. Online misinformation is linked to early COVID-19 vaccination hesitancy and refusal. *Scientific Reports* 12, 1 (2022), 5966.
- [91] Moritz Pilarski, Kirill Olegovich Solovev, and Nicolas Pröllochs. 2024. Community Notes vs. Snoping: How the crowd selects fact-checking targets on social media. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 18. 1262–1275.
- [92] Ethan Porter and Thomas J Wood. 2021. The global effectiveness of fact-checking: Evidence from simultaneous experiments in Argentina, Nigeria, South Africa, and the United Kingdom. *Proceedings of the National Academy of Sciences* 118, 37 (2021), e2104235118.
- [93] Nicolas Pröllochs, Dominik Bär, and Stefan Feuerriegel. 2021. Emotions explain differences in the diffusion of true vs. false social media rumors. *Scientific Reports* 11, 1 (2021), 22721.
- [94] Nicolas Pröllochs. 2022. Community-based fact-checking on Twitter’s Birdwatch platform. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16. 794–805.
- [95] Steve Rathje, Jay J. Van Bavel, and Sander van der Linden. 2021. Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences* 118, 26 (2021), e2024292118.
- [96] Raquel Recuero, Felipe Bonow Soares, and Anatoliy Grudz. 2020. Hyperpartisanship, disinformation and political conversations on Twitter: The Brazilian presidential election of 2018. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 569–578.
- [97] Thomas Renaud, David Restrepo Amariles, and Aurore Troussel. 2024. Collaboratively adding context to social media posts reduces the sharing of false news. *ArXiv* (2024).
- [98] Manoel Horta Ribeiro, Robert West, Ryan Lewis, and Sanjay Kairam. 2025. Post guidance for online communities. *Proceedings of the ACM on Human-Computer Interaction* forthcoming (2025).
- [99] Claire E Robertson, Nicolas Pröllochs, Kaoru Schwarzenegger, Philip Pärnamets, Jay J Van Bavel, and Stefan Feuerriegel. 2023. Negativity drives online news consumption. *Nature Human Behaviour* 7, 5 (2023), 812–822.
- [100] Hannes Rosenbusch, Anthony M Evans, and Marcel Zeelenberg. 2019. Multilevel emotion transfer on YouTube: Disentangling the effects of emotional contagion and homophily on video audiences. *Social Psychological and Personality Science* 10, 8 (2019), 1028–1035.
- [101] Katja Rost, Lea Stahel, and Bruno S Frey. 2016. Digital social norm enforcement: Online firestorms in social media. *PLoS ONE* 11, 6 (2016), e0155923.
- [102] David Rozado, Ruth Hughes, and Jamin Halberstadt. 2022. Longitudinal analysis of sentiment and emotion in news media headlines using automated labelling with Transformer language models. *PLoS ONE* 17, 10 (2022), e0276367.
- [103] Mohammed Saeed, Nicolas Traub, Maelle Nicolas, Gianluca Demartini, and Paolo Papotti. 2022. Crowdsourced fact-checking at Twitter: How does the crowd compare with experts?. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 1736–1746.
- [104] Jessica M Salerno and Liana C Peter-Hagene. 2013. The interactive effect of anger and disgust on moral outrage and judgments. *Psychological Science* 24, 10 (2013), 2069–2078.
- [105] Yoonmo Sang, Jee Young Lee, Sora Park, Caroline Fisher, and Glen Fuller. 2020. Signalling and expressive interaction: Online news users’ different modes of interaction on digital platforms. *Digital Journalism* 8, 4 (2020), 467–485.
- [106] Disa A Sauter, Frank Eisner, Paul Ekman, and Sophie K Scott. 2010. Cross-cultural recognition of basic emotions through nonverbal emotional vocalizations. *Proceedings of the National Academy of Sciences* 107, 6 (2010), 2408–2412.
- [107] Brennan Schaffner, Arjun Nitin Bhagoji, Siyuan Cheng, Jacqueline Mei, Jay L Shen, Grace Wang, Marshini Chetty, Nick Feamster, Genevieve Lakier, and Chenhao Tan. 2024. “Community guidelines make this the best party on the internet”: An in-depth study of online platforms’ content moderation policies. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [108] Zakir Shah, Jianxun Chu, Sara Qaisar, Zameer Hassan, and Usman Ghani. 2020. Perceived public condemnation and avoidance intentions: The mediating role of moral outrage. *Journal of Public Affairs* 20, 1 (2020), e2027.
- [109] Leah M Smith, Linda E Lévesque, Jay S Kaufman, and Erin C Strumpf. 2017. Strategies for evaluating the assumptions of the regression discontinuity design: A case study using a human papillomavirus vaccination programme. *International Journal of Epidemiology* 46, 3 (2017), 939–949.
- [110] Kirill Solovev and Nicolas Pröllochs. 2023. Moralized language predicts hate speech on social media. *PNAS Nexus* 2, 1 (2023), pgae281.
- [111] Kirill Solovev and Nicolas Pröllochs. 2022. Moral emotions shape the virality of COVID-19 misinformation on social media. In *Proceedings of the ACM Web Conference 2022*. 3706–3717.
- [112] SSRS. 2023. CNN Poll on Biden, economy and elections. <https://www.documentcloud.org/documents/23895856-cnn-poll-on-biden-economy-and-elections>.

- [113] Statista. 2024. Number of internet and social media users worldwide as of July 2024. <https://www.statista.com/statistics/617136/digital-population-worldwide/>.
- [114] Stefan Stieglitz and Linh Dang-Xuan. 2013. Emotions and information diffusion in social media—sentiment of microblogs and sharing behavior. *Journal of Management Information Systems* 29, 4 (2013), 217–248.
- [115] Wienke Strathern, Mirco Schoenfeld, Raji Ghawi, and Juergen Pfeffer. 2020. Against the others! Detecting moral outrage in social media networks. In *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. 322–326.
- [116] Huiyun Tang, Gabriele Lenzini, Samuel Greiff, Björn Rohles, and Anastasia Sergeeva. 2024. “Who knows? Maybe it really works”: Analysing users’ perceptions of health misinformation on social media. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 1499–1517.
- [117] Jay J Van Bavel, Claire E Robertson, Kareena Del Rosario, Jesper Rasmussen, and Steve Rathje. 2024. Social media and morality. *Annual Review of Psychology* 75, 1 (2024), 311–340.
- [118] Julie Van de Vyver and Dominic Abrams. 2015. Testing the prosocial effectiveness of the prototypical moral emotions: Elevation increases benevolent behaviors and outrage increases justice behaviors. *Journal of Experimental Social Psychology* 58 (2015), 23–33.
- [119] Frances van der Eijk and Simon Columbus. 2023. Expressions of moral disgust reflect both disgust and anger. *Cognition and Emotion* 37, 3 (2023), 499–514.
- [120] Rama Adithya Varanasi, Joyojeet Pal, and Aditya Vashistha. 2022. Accost, accede, or amplify: Attitudes towards COVID-19 misinformation on WhatsApp in India. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [121] Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151.
- [122] Senuri Wijenayake, Danula Hettiachchi, Simo Hosio, Vassilis Kostakos, and Jorge Goncalves. 2020. Effect of conformity on perceived trustworthiness of news in social media. *IEEE Internet Computing* 25, 1 (2020), 12–19.
- [123] Tamar Wilner, Kayo Mimizuoka, Ayesha Bhimdiwala, Jason C Young, and Ahmer Arif. 2023. It’s about time: Attending to temporality in misinformation interventions. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [124] Stefan Wojcik, Sophie Hilgard, Nick Judd, Delia Mocanu, Stephen Ragain, MB Hunzaker, Keith Coleman, and Jay Baxter. 2022. Birdwatch: Crowd wisdom and bridging algorithms can inform understanding and reduce the spread of misinformation. *ArXiv* (2022).
- [125] X. 2024. About replies and mentions. <https://help.x.com/en/using-x/mentions-and-replies>
- [126] Ali Yadollahi, Ameneh Gholipour Shahraki, and Osmar R Zaiane. 2017. Current state of text sentiment analysis from opinion to emotion mining. *ACM Computing Surveys (CSUR)* 50, 2 (2017), 1–33.
- [127] Waheeb Yaqub, Otari Kakhidze, Morgan L Brockman, Nasir Memon, and Sameer Patil. 2020. Effects of credibility indicators on social media news sharing intent. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [128] YouTube. 2024. Testing new ways to offer viewers more context and information on videos. <https://blog.youtube/news-and-events/new-ways-to-offer-viewers-more-context/>
- [129] Himanshu Zade, Spencer Williams, Theresa T Tran, Christina Smith, Sukrit Venkatagiri, Gary Hsieh, and Kate Starbird. 2024. To reply or to quote: Comparing conversational framing strategies on Twitter. *ACM Journal on Computing and Sustainable Societies* 2, 1 (2024), 1–27.
- [130] Yixuan Zhang, Joseph D Gaggiano, Nutchanon Yongsatianchot, Nurul M Suhaimi, Miso Kim, Yifan Sun, Jacqueline Griffin, and Andrea G Parker. 2023. What do we mean when we talk about trust in social media? A systematic review. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–22.
- [131] Yixuan Zhang, Yimeng Wang, Nutchanon Yongsatianchot, Joseph D Gaggiano, Nurul M Suhaimi, Anne Okrah, Miso Kim, Jacqueline Griffin, and Andrea G Parker. 2024. Profiling the dynamics of trust & distrust in social media: A survey study. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–24.
- [132] Jiawei Zhou, Yixuan Zhang, Qianni Luo, Andrea G Parker, and Munmun De Choudhury. 2023. Synthetic lies: Understanding AI-generated misinformation and evaluating algorithmic and human solutions. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [133] Fabiana Zollo, Petra Kralj Novak, Michela Del Vicario, Alessandro Bessi, Igor Mozetič, Antonio Scala, Guido Caldarelli, and Walter Quattrociocchi. 2015. Emotional dynamics in the age of misinformation. *PLoS ONE* 10, 9 (2015), e0138740.