Overestimation in the Aggregation of Emotional Intensity of Social Media Content

Accepted by the *Journal of Personality and Social Psychology: Attitudes and Social Cognition* on May 12, 2025.

Jonas P. Schöne[1, 2, *], Matthew D. Rocklage[3], Brian Parkinson[4], Amit Goldenberg[5, 6, 7]

Authors Affiliations:

[1] Human-Centered AI, Stanford University

[2] Department of Psychology, Stanford University

[3] D'Amore-McKim School of Business, Northeastern University

[4] Department of Experimental Psychology, University of Oxford

[5] Harvard Business School, Harvard University

[6] Harvard Department of Psychology, Harvard University

[7] Digital, Data and Design Institute, Harvard University

* Corresponding Author

Email: jschoene@stanford.edu

Data availability

The data and analysis for Studies 1–5 is available at https://osf.io/y2pj5/?view_only=9cb21b1d5c6c4b38824d15405abfd6f1.

Code availability

The code for the analysis of Studies 1–5 is available at https://osf.io/y2pj5/?view_only=9cb21b1d5c6c4b38824d15405abfd6f1. The code for the tasks can be found at https://github.com/GoldenbergLab/task-sequential-word-text-estimation.

Author Contributions: J.P.S.: designed research, performed research, contributed new reagents/analytic tools, analyzed data, and wrote the paper. M.D.R.: contributed new reagents/analytic tools and analyzed data. B.P.: designed research and wrote the paper. A.G.: designed research, contributed new reagents/analytic tools, analyzed data, and wrote the paper.

Competing Interest Statement: The authors had no conflict of interest to declare.

Abstract

Users on social media are regularly presented with sequences of emotional content in their newsfeeds, which affects their viewpoints and emotions. Could the way users aggregate and remember emotional content from their feeds contribute to the fact emotions are amplified on social platforms? Across five studies (N = 1,051), using experimentally manipulated social media feeds, we found that participants consistently overestimated the average emotional intensity of the individual responses expressed by other users in a sequence (Study 1a). This overestimation led to stronger emotional reactions to the news content that these responses were reacting to (Study 1b). Investigating the mechanism suggested that while there was stronger memory for more emotional responses within a response sequence, we could not find a direct link between memory and overestimation (Study 2). We showed that overestimation was driven mainly by the salience of emotional intensity of different items in the sequence, by replicating the effect using sequences of emotional words (Study 3). We then turned to the consequences of overestimation, showing that overestimation of emotional sequences was uniquely associated with perceiving more intense emotional responses as more representative of how other people would react (Study 4), and with overestimation of the emotionality of the newsfeed as a whole (Study 5). Overestimation of the average individual emotional intensity ratings of a sequence was also predictive of willingness to share articles. This set of findings sheds light on how sampling from newsfeeds amplifies the perception of emotionality.

Keywords: Emotion, Perception, Social Media, Sequential Presentation, Emotional Norms

Statement of Limitations

Several limitations may affect the generalizability of our findings. First, the external validity is constrained because participants were exposed to fictional news articles without background information about the content creators. In real social media settings, prior knowledge of the content creators can influence users' perceptions, potentially reducing the overestimation observed in our study. Second, the number of responses to each news article was randomized, unlike real social media, where response volume can signal public interest and emotional intensity. Third, the random selection of responses may not reflect real-world scenarios, where emotional responses are often more clustered and aligned with the emotional intensity of the content. This could lead to a more accurate estimation of average emotionality than observed in our study. Fourth, the study measured hypothetical willingness to share, leaving open the question of real-world sharing behavior. Fifth, because all events were fictional, we could not assess how overestimation of emotions might affect participants' real beliefs. Finally, responses always matched the valence of the news article, while in real-world discussions, comments often express disagreement or alternative perspectives. This controlled design helped isolate the role of emotional intensity, but it may not fully capture the complexity of online emotional expression.

The average American scrolls through 300 feet (91.44 m) of online content per day (Wade, 2017). Unsurprisingly, this amount of content plays a crucial role in shaping their perceptions and feelings (Kelly & Sharot, 2023). An important component of newsfeeds is the emotions expressed in users' responses (Brady, McLoughlin, et al., 2023; Goldenberg & Gross, 2020). Expressed emotions provide a window on others' internal states, and people are both highly tuned to others' emotions as well as adept at deciphering emotions swiftly and effectively (Eimer & Holmes, 2007; Schirmer & Adolphs, 2017).

Users' understanding of their feeds requires aggregating multiple emotional responses into overall representations of what others think and feel. How do users generate this aggregated representation of emotionality? Do they weigh all information equally or are they unduly influenced by responses that are more salient or more easily remembered? Answering these questions is important because people's behavior on social media is contingent on their understanding of the norms within the platform. A systematic overestimation of aggregated emotional content could distort users' judgments of what others feel and influence users' subsequent engagement with the platform (see overview Figure 1).

Emotions on Social Media

Social media is an attention economy where both platforms and users are incentivized to create content that draws attention. One way to capture more attention is through expressing stronger emotions, which tends to result in more engagement for emotional content (Brady et al., 2019; Goldenberg & Gross, 2020). As social media platforms want users to be engaged with the platform, they designed algorithms that prioritize recommending it to wider audiences (Brady, Jackson, et al., 2023), further increasing users' exposure to emotional content. Beyond this top-down algorithmic influence, users are also motivated to produce more emotionally intense content themselves because they are sensitive to the rewarding feedback they receive when others engage with their emotional posts (Brady et al., 2021; Lindström et al., 2021). As a result of algorithmic reinforcement and user behavior, emotional content is very prevalent across social media newsfeeds.

Given the central role of emotional content in our online experiences and the dissemination of news, previous research has examined how such content is perceived and shared. For instance, Brady et al. (2023) investigated whether the emotional intensity intended by content creators aligns with how it is perceived by observers. Their findings suggested that observers, particularly frequent social media users, tended to overestimate the emotional intensity of a single post compared to creators' self-reports of the emotional intensity they intended to convey. In terms of how emotionality affects sharing behavior, studies consistently show that greater emotional intensity leads to more sharing (Robertson et al., 2023; J. Schöne et al., 2021; J. P. Schöne et al., 2023; Zhang & Qu, 2020).

Much of the work on emotionality on social media has focused on how people react and respond to a single post. However, one unique aspect of social media feeds is that emotional reactions to news and events are often numerous. In evaluating the emotional reaction to a certain situation, users often need to aggregate emotional information generated by multiple responses. How do users' aggregate sequences of emotional responses, and are they accurate in doing so?

Overestimation in the Aggregation of Emotions

We argue that when exposed to a stream of messages, users will overestimate the emotional intensity of others. This prediction is similar to those made by the peak-end rule, which posits that a series of emotional experiences are evaluated by averaging the most intense and the last moments of the sequence (Redelmeier & Kahneman, 1996). Unlike this prediction, which emphasizes only the highest and last experiences (see SI for comparison), we argue that the potential mechanism is that people are more likely to remember higher-intensity content as it is more salient (Gorges et al., 2024; Jackson, 2018; Jackson et al., 2009). The overestimation is likely to be more pronounced in longer response sequences, as they are more likely to contain a higher number of high-intensity stimuli. Finally, since people often pay more attention to negative emotions (Fazio et al., 2015; Rozin & Royzman, 2001), they may overestimate the estimated average emotional intensity of negative sequences to a greater extent than positive sequences.

Although no prior studies have directly investigated the aggregation of emotional content in social media feeds, indirect evidence for intensity overestimation can be found in a recent paper that examined how people aggregate emotional information in sequences of faces (Goldenberg et al., 2022). Participants were shown a series of facial expressions asked to evaluate the average emotional intensity of the sequence. Results suggested that participants overestimated the estimated average emotional intensity of the sequence, compared to the average perceived intensity of each individual face. One potential explanation for this overestimation lies in the role of emotional salience in memory and judgment. Emotionally intense stimuli tend to be more attention-grabbing (Öhman et al., 2001) and are more likely to be remembered (for an overview, Kensinger & Schacter, 2008). When individuals process multiple emotional stimuli in sequence, the most intense items stand out and become more accessible in memory. Consequently, when participants later estimate the emotional intensity of the entire sequence, their judgments are disproportionately influenced by these more salient, memorable moments, leading to an overall overestimation.

This initial evidence for overestimation in face sequences, along with findings from emotional memory research supporting the role of salience, raises the question of whether the same mechanism applies to averaging textual content. The perceptual averaging of visual cues, such as those provided by faces, operates in the same manner as the aggregation of semantic stimuli, such as those contained in sequences of text. While processing of faces requires automatic visual processing, textually presented affect requires both visual and semantic processing to decode word meaning (Derks et al., 2008; Harris & Paradice, 2007). It is therefore possible that findings from the processing of visual cues do not translate to textual stimuli. Given the importance of understanding how users aggregate and represent emotional information in sequences of written content, further investigation is required of how this kind of material is processed and combined.

Overestimation of the average emotional intensity of social media content could lead to significant consequences for users' perception and behavior on social media. First, inaccurate perceptions of others' responses may amplify users' own emotional responses to the associated events (Manstead & Fischer, 2001; Parkinson & Simons, 2009). Second, users may incorrectly infer that the normative response is more intense than it actually is. This, in turn, might influence which content users decide to share and produce (Brady et al., 2019; Frimer et al., 2023; Schöne et al., 2021). People naturally tend to share emotional content, a phenomenon known as the social sharing of emotions (Luminet IV et al., 2000). For instance, during emotionally charged events such as the COVID-19 pandemic, individuals created and shared more related content on social media (Ma et al., 2024). Prior research on social sharing suggests that people are particularly inclined to share emotionally intense (Rimé, 2009), with a stronger tendency to share negative content, as it is perceived as more urgent or socially relevant (Robertson et al., 2023). If users systematically overestimate the emotionality of a sequence, this effect may further increase their willingness to share the corresponding news article. Third, users may arrive at the general impression that social media content as a whole is more emotionally intense than it actually is. Providing some initial support to this last prediction is the fact that users who use social media frequently perceive a greater level of outrage on these platforms than those who spend less time online (Brady, McLoughlin, et al., 2023).

The Current Research

In the present research, we conducted six studies to test whether people overestimate the average emotional intensity of sequences of feed-like text responses and sought to identify underlying mechanisms and consequences. Studies 1a and 1b aimed to establish whether participants overestimated the emotional intensity of sequences of text presented as responses to fictional news articles. The three main pre-registered hypotheses were: that people would overestimate the intensity of the average expressed emotion (H1); that overestimation would be stronger for longer sequences (H2); and for sequences of negative emotion expressions (H3). In Studies 2 and 3, we wanted to understand the mechanisms driving the overestimation of emotional sequences. In Study 2, we tested whether participants were more likely to remember strong expressions of emotions. Study 3 tested whether salience alone drives the effect by presenting participants with sequences of individual emotional words instead of full text passages, assessing whether overestimation occurs without contextual information. Studies 4 and 5 were designed to examine the potential consequences of overestimation. Study 4 assessed whether overestimation of sequences was associated with the perception that more emotionally intense responses are more representative of the general public's response. Finally, Study 5 investigated whether individuals tended to overestimate their newsfeed as a whole, rather than specific sequences.

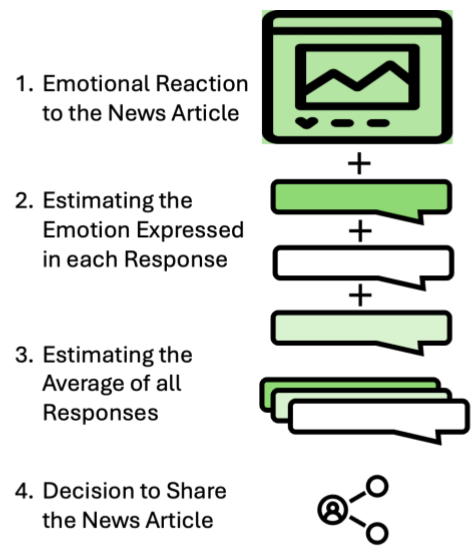

Figure 1. Schematic representation of sequential emotional estimations on social media platforms. Users initially experience an emotional reaction to the content they encounter (1). To grasp how others feel, they integrate the individual emotions expressed by other users (2) into an aggregate value (3). Finally, based on the user's evaluation of the emotionality of the responses and their emotional reaction to the news article they may decide to share it with their network (4).

Transparency and Openness

For each of the five studies, we report how we determined our sample size, all data exclusions (if any), all manipulations, and all measures used in the study or the Supplementary Material). All data, analysis code, pre-registrations, and research materials are available at:

Overview: https://osf.io/y2pj5/?view_only=9cb21b1d5c6c4b38824d15405abfd6f1

Study 1a: https://osf.io/y2pj5/?view_only=24137366f510428da1804cafd6ac1a36

Study 1b: https://osf.io/y2pj5/?view_only=fcc3c57b21c84f21ae79f833757e5198

Study 2: https://osf.io/y2pj5/?view_only=4d98688f64ea49e2bc8cdd78008f9f51

Study 3: https://osf.io/yj8ku/?view_only=51a69939ab7842f395feef51620b8828

Study 4: https://osf.io/y2pj5/?view_only=08a24e89593a4863a47d86db51841f5b

Study 5: https://osf.io/y2pj5/?view_only=661f911e84af466588a80bbcba8d9934

For all subsequent linear mixed modeling analyses, we used the lme4 package (version 1.1.2; Bates et al., 2013) and the lmerTest package (version 3.1.2; (Kuznetsova et al., 2013), both implemented in RStudio (version 4.3.2).

Study 1a-b: Overestimation of the emotional intensity of sequences of text.

The goal of Studies 1a and 1b was to determine whether participants overestimated the average emotional intensity of individual responses in a sequence following a news article. In Study 1a, the news article was presented before the response sequence to examine whether the emotional reaction to the news article influenced the extent of overestimation of the response sequence. In Study 1b, the order was reversed, with the sequence of responses presented first, to assess whether the degree of overestimation influenced their emotional reaction to the news article. Study 1b also allowed us to examine whether participants' willingness to share the news article was affected by the extent of their overestimation of the response sequence.

Method

Ethics statement

All studies, including Studies 1a and 1b, were approved by the Human Subject Committee at two institutions (*redacted to peer-review*). All participants provided informed consent and received financial compensation for their participation.

Stimuli

Generating Fictional News Articles. Before conducting the main studies, we needed to create fictional news articles and responses to them. First, we manually developed and wrote 50 stories, comprising both positive (25) and negative stories (25) in a journalistic style similar to typical articles in local newspapers. These stories covered a wide range of topics, including sports, medicine, weather, and local news, drawing inspiration from real news articles. Each story consisted of a title (e.g. "*Anonymous Donor Gifts Entire Hospital Wing"*), a concise 2 to 3-sentence summary, and an accompanying image (Figure 2, see SI for all stories and responses). Importantly, none of these stories described actual events or featured real individuals and also avoided covering political topics or societal issues for which people might hold prior opinions and emotional reactions. Our objective was to create an equal number of articles eliciting positive and negative emotions at varying intensities. We chose not to elicit specific distinct emotions, because genuine social media articles often involve a mix of emotions. Additionally, we aimed to ensure that the emotional intensity of both negative and positive news articles was similar (for additional details, please see the Materials and Methods section; the full text if all the news stories are included in the Supplementary Information).

Figure 2. Examples of fictional news articles and user responses**.** News articles and their responses are categorized as either positive (shown on the left) or negative (shown on the right). Each news article consists of a title, an image, and a short description summarizing its content (top images). The responses are plain text without any further information about the content producer (bottom images).

Generating User Responses to Articles. After creating a set of news articles, our goal was to assess the intensity of participants' emotional reactions to the articles and generate human text responses for each one (see Materials and Methods for more details). To accomplish this, we conducted a pilot study in which participants (41 females, 41 males, age: M = 26.30, SD = 8.60, range = 18-71) first rated their *own emotional reactions* to a sample of the 50 news stories on a scale from 0 (neutral) to 9 (highly emotional) by asking, “*please rate the emotionality of the response to the news story.*” Each participant rated 30 news articles. This resulted in an average of 49.20 ratings per news article. Using a t-test to compare the emotional reaction scores between positive and negative news stories, we found no significant difference between the positive (M = 5.30, SD = 0.92) and negative stories (M = 5.50, SD = 0.87) in terms of people's emotional reaction intensity (t ((47.90) = -0.94), p = .35).

We also used this pilot to gather text responses to the news articles. After rating each article, participants were asked to generate a text response as if they were posting on social media. They were instructed to write one or two sentences, with a minimum of twelve characters. The instructions read: "*Please create a text response that expresses an emotional intensity of 3 (0: Not at all Emotional to 9: Very Emotional). The length of your response should be one sentence long (at least twelve characters)*". To help participants understand how to adjust their responses based on emotional intensity, they were provided with two example responses of varying intensities during the two practice trials. To ensure a range of emotional intensities, each participant was randomly assigned an emotional intensity level from 0 (neutral) to 9 (very high emotional intensity). They were then asked to write their response according to this assigned intensity, ensuring that their emotional expression reflected the level they were given. In total, participants created 2,486 responses to 50 stories, averaging approximately 30 responses per story. We excluded responses using these criteria: exceeding 120 characters, not fitting the context of the news article (e.g. if they misunderstood the article to the extent that their response did not logically correspond to it), containing profanities, difficult to understand, using sarcasm or humor. We also made sure that the valence of the news articles and responses matched, meaning the responses spanned a range of intensities from neutral to negative or neutral to positive, according to the valence of the news article.

For Study 1a, we selected 20 responses per article, ensuring that each set included two responses for each level of the pre-assigned emotional intensity scale (0–9), matching the positive or negative valence of the article. However, these pre-assigned intensity levels were not used to calculate the individual sequence average in the studies. Instead, we relied on participants' own individual ratings to determine the perceived emotional intensity. Utilizing the average emotional intensity ratings of the responses obtained from Study 1a, we conducted a retrospective analysis to examine the correlation between the intended intensity levels and the actual ratings provided during the experiment. The analysis revealed a significant correlation between the intended intensity levels, determined by selecting the stimuli based on the pre-assigned intensity of the pilot, and the average intensity ratings from Study 1a ($r = .79$, $t(998) = 40.94$, $p < .001$). These results suggest a close alignment between the intended intensities and the subsequently rated intensities for the selected stimuli.

For Study 1b and subsequent studies, we refined the selection of news articles as well as their responses using participants' ratings from Study 1a to ensure that positive and negative content were comparable in both distribution and emotional intensity. This selection process aimed to ensure that any observed effects in Study 1a were not driven by pre-existing intensity differences between positive and negative content but rather reflected participants' perceived emotional intensity. We first chose 40 news articles (20 positive and 20 negative) such that the mean emotional reaction intensity of the positive and negative stories were not significantly different (positive: M = 63.24, SD = 10.7; negative: M = 64.12, SD = 13.5, t (36.15) = 0.22, p = .79). Next, for each of the 40 articles, we selected 15 responses based on two criteria: (1) The mean individual emotional intensity rating of all positive responses (M = 53.09, SD = 17.01) was not significantly different from the mean individual emotional intensity rating of the negative responses (M = 54.79, SD = 15.20, t (590.15) = 1.29, p = .19); and (2) the emotional reaction in response to the positive (W = 0.90, p = .053) and negative (W = 0.97, p = .85) news articles was normally distributed according to Shapiro-Wilk tests.

In summary, Study 1a utilized 50 news articles with 20 responses each, while Study 1b used 40 news articles with 15 responses each, matched according to the emotional intensity ratings collected in Study 1a.

Participants

For Studies 1a and 1b, as well as all other studies, we limited our recruitment to individuals residing in the United Kingdom or the United States to ensure that participants were fluent in reading English and could accurately judge their emotional intensity. For Study 1a, based on our power analysis, we concluded that 300 participants would be needed to achieve 80% power in detecting an overestimation effect, assuming a similar effect size to previous studies (see SI). We therefore recruited 300 native English speakers from the United Kingdom and the United States via Prolific. Following our pre-registered exclusion criteria, we excluded 11 participants who failed more than three attention checks or missed survey responses. In line with our pre-registered criteria, we excluded 7 participants whose mean estimation exceeded 2.5 standard deviations from the total sample. We chose this exclusion criterion for all the studies because selecting only the extreme values, whether low or high, likely indicates that ratings were conducted quickly to finish the task, without regard to the averages. Such extreme values could further skew the statistical analysis. This resulted in a final sample size of 282 participants (145 females, 135 males, 1 other, 1 who prefers not to say, age: M = 25.35, SD = 7.37, range = 18-70). Participants were compensated $4.60 for their participation in the 40-minute task.

We used the effect size of Study 1a to estimate the number of participants for Study 1b. Our power analysis suggested to detect a similar effect as in Study1a with an 80% probability, we would need 60 participants completing 15 trials (see SI). However, due to uncertainty about whether the degree of overestimation of emotions would be smaller when presenting the news article at the end of the trial, we opted to recruit a larger sample than suggested by the power analysis. Specifically, we recruited 150 native English speakers from the United Kingdom and the

United States via Prolific, with each participant completing 15 trials. We excluded 8 participants who failed more than 3 attention checks or had missing data. Additionally, following the same pre-registered exclusion criteria as in Study 1a, we excluded 2 more participants whose mean estimation was more than 2.5 standard deviations away from the sample mean. The final sample consisted of 139 participants (70 females, 67 males, 1 other, 1 who preferred not to say, age: M = 24.68, SD = 6.79, range = 18-57). In Study 1b, participants were compensated $4.60 for their participation in the 40-minute task as in Study 1a.

Procedure

Sequential responses task. Our task involved 15 trials preceded by an additional two practice trials in both studies, which were excluded from the analysis. At the beginning of each trial (see Figure 3), participants were randomly presented with one of 50 fictional news articles. After three seconds, a rating scale appeared, prompting them to rate their own emotional reaction on a scale from 0 (Not at all Emotional) to 100 (Very Emotional) by asking, "*Please rate your emotions in response to the news story.*" Instead of using the 0 – 9 Likert scale used in the stimulus evaluation pilot, we switched to a 0 – 100 slider in Study1a to allow for more fine-grained variation and to make it more difficult for a participant to remember the exact rating that they have given to each stimulus. Next, participants saw a fixation cross in the center of the screen for a randomly determined duration of 400-600ms, after which one of the 20 possible responses was presented to the participant. After 1 second a scale appeared, and participants were asked to rate the individual emotional intensity of the response on the same 1 to 100 scale with the phrasing, "*Please rate the emotionality of the response to the news story*". Using individual ratings for each response allowed us to obtain a more precise estimation of the actual mean without relying on scores from previous raters as other studies have done. The number of responses presented in each trial varied randomly between 4 and 10, with each response preceded by a fixation cross and followed by the participant's rating. The sequence length for each trial was determined at random, ensuring variation in the number of responses participants encountered across trials. We chose a minimum sequence length of 4 because previous studies indicated that overestimation occurs at this length (Goldenberg et al., 2022). After viewing and rating all the responses presented in a trial, participants were asked to estimate the average emotional intensity of the sequence using the same scale by asking *"Please rate the average emotionality of the sequence of responses to the news story."* We used the term "average" to encourage participants to make a deliberate numerical estimation, aligning more closely with how we calculated the actual mean. To measure overestimation, we compared the estimated average emotional intensity to the average calculated by using the individual ratings provided by each participant (actual mean). Each news article and response were presented only once to each participant. After completing the main task, participants filled out a short survey. The survey for all studies included demographic questions such as age, gender, and race, as well as individual differences assessed by several scales. We included the Social Interaction Anxiety (SIAS-6, Mattick & Clarke, 1998) based on previous findings that individuals with higher social interaction anxiety show greater overestimation of emotional expressions in face sequence and crowds (Goldenberg et al., 2021, 2022), suggesting that some people may be more sensitive to others' emotional expressions, which could influence how they perceive and aggregate emotional content. Because emotional reactivity and personality traits may also shape how individuals process and integrate emotional information, we assessed emotional reactivity using the PERS Reactivity Scale (Becerra et al., 2019) and measured personality traits, particularly neuroticism, using the Big-Five Personality Scale (Gosling et al., 2003). These measures allowed us to explore whether individual differences in emotional responsiveness and personality influence overestimation effects. (See SI for differences in overestimation of the emotional intensity of the sequence with survey items).

Study 1b used the same task structure as the previous experiment, including 2 practice trials and 15 actual trials, but with three modifications. First, unlike the previous experiment, participants in Study 1b first saw and rated individual responses and then estimated their average emotional intensity of these responses in the sequence to the news article before seeing the news article itself. Each trial began with a series of 4–10 responses, where participants rated the emotional intensity of each response one at a time. Each rating was followed by a fixation cross before the next response appeared. After rating all responses in the sequence, participants were then asked to rate the estimated average emotional intensity of the sequence using the same 0–100 scale. Only after completing this rating did they finally see the corresponding news article, at which point they rated their emotional reaction to it. Second, we selected 40 out of the original 50 news articles and 15 responses from the original set of 20 responses each to ensure that the emotional reaction to the news article and the emotional intensity of the responses were identical for positive and negative sequences. Finally, participants were asked to indicate the likelihood of sharing the news story on social media using a scale ranging from 0 (very unlikely) to 100 (very likely) at the end of each trial. After completing the primary task, participants were directed to an identical brief survey as in Study 1a (see SI).

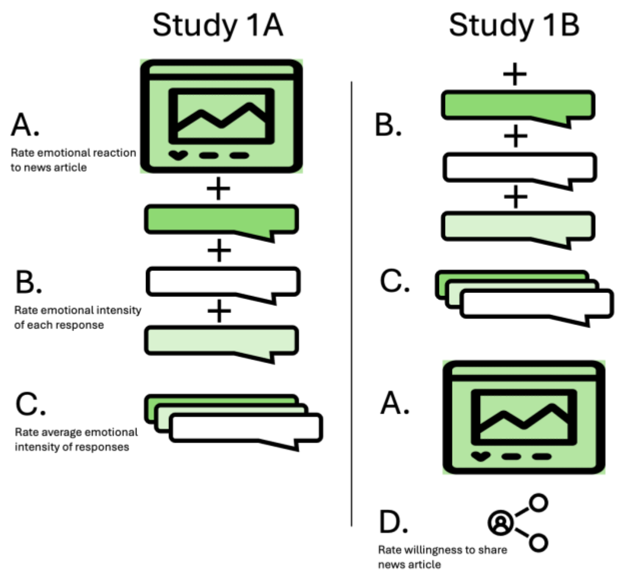

Figure 3. Design for Study 1a and Study 1b. In both studies, participants rated their emotional reaction to the news article (A). After a brief fixation cross, the emotional intensity of each response in isolation (B), and the estimated average emotional intensity of the sequence (C). In Study 1a, the rating of their emotional reaction to the news article (A) was conducted at the beginning, whereas in Study 1b, it was conducted at the end. Additionally, in Study 1b, participants rated their hypothetical willingness to share the news article (D).

Results

Evaluating Overestimation of the Response Sequences. To evaluate whether participants overestimated the average emotional intensity of response sequences, we compared the actual average of their ratings of individual responses (rated immediately after reading each response), to their evaluation of the estimated average emotionality of the entire response sequence. Using a mixed model repeated measures analysis, we tested whether participants' evaluations of the estimated average emotional intensity of the sequence exceeded the average of their individual ratings, as predicted by our hypothesis (H1). Specifically, the model compared the rating values of two rating types: the individual average rating and the estimated average emotional intensity of the sequence for each trial. The rating type was included as a dummy categorical variable in the model, with the individual average rating as the base level for comparison (see SI for detailed model description). In Study 1a, we included the emotional reaction intensity of the news articles as a covariate to determine whether there was an overestimation effect beyond the impact of the news article. In Study 1b, since the news article was introduced only at the end of each trial after the sequence of responses had been rated already, we did not need to include its emotional reaction intensity as a covariate. Additionally, we incorporated random intercepts for participants and news articles to address individual differences in emotional intensity perception and news article reaction intensities, respectively.

We found that participants overestimated the emotional intensity of the sequence compared to the average of their own ratings of the individual responses in the sequence (H1: Study 1a, *b* = 1.86, *SE* = 0.29, *t* (8121.17) = 6.21, $p < .001$, $R^2 = .11$, 95% Confidence Intervals = [1.27, 2.45]; Study 1b, *b* = 2.10 [1.35, 2.86], *SE* = 0.38, *t* (3968.97) = 5.46, $p < .001$, $R^2 = .12$).

To examine how sequence length (H2) and valence (H3) influenced participants' tendency to overestimate the average individual emotion expressed in the response sequence, we created a difference score between participants' estimation of the average responses and the average of the individual emotion intensity ratings of the responses in a trial. We then used sequence length and its interaction with valence of each trial as predictors of this difference score, including random intercepts for participants and news stories (and the covariate of news article intensity for Study 1a only, see SI for detailed model description).

Results indicated that overestimation was more pronounced for longer sequences (H2: Study 1a, b = 0.39, *SE* = 0.12, *t* (4071.48) = 3.21, $p < .001$, $R^2 = .03$, 95% Confidence Intervals = [0.15, 0.63]; Study 1b, *b* = 0.61 [0.40, 0.86], *SE* = 0.17, *t* (2030.00) = 3.65, $p < .001$, $R^2 = .01$, see Figure 4). However, there was no difference in estimations of positive and negative sequences (H3: Study 1a, *b* = -0.62, *SE* = 1.31, *t* (1025.43) = -0.46, $p = .63$, $R^2 = .03$, 95% Confidence Intervals = [-3.20, 1.97], Study 1b, *b* = -0.25 [-3.72, 3.22], *SE* = 0.61, *t* (1228.86) = -0.14, $p = .87$, $R^2 = .01$). Additionally, we tested whether trial position had a similar cumulative influence as sequence length by repeating the analysis with trial position as the predictor. We did not find a significant effect, and the relevant analyses are presented in the Supplementary Information.

Overestimation of the Response Sequence and Emotional Reaction to the News Article. We examined the relationship between participants' emotional reaction to the news article and the overestimation of the sequence in Study 1a, and the reverse relationship in Study 1b. Analyzing the influence of the intensity of the reaction to the news article in Study 1a using the same model as for H 2-3, we found a significant positive effect on overestimation: participants overestimated the intensity of the response sequence to a greater extent when the news article was rated as more emotionally intense (*b* = 0.065, *SE* = 0.0073, *t* (1836.57) = 8.86, $p < .001$, $R^2 = .03$, 95% Confidence Intervals = [0.05, 0.08]). In Study 1b, we modified the model used in Study 1a by switching the dependent variable to estimation difference and the independent variable to news article reaction rating. We found a positive association between the extent of participants' overestimation of the average individual emotional intensity expressed in the response sequences and the intensity of their emotional reaction to the news article (b = 0.19, SE = 0.040, t (1964.00) = 4.67, p < .001, R^2 = .007, 95% Confidence Intervals = [0.11 – 0.27]).

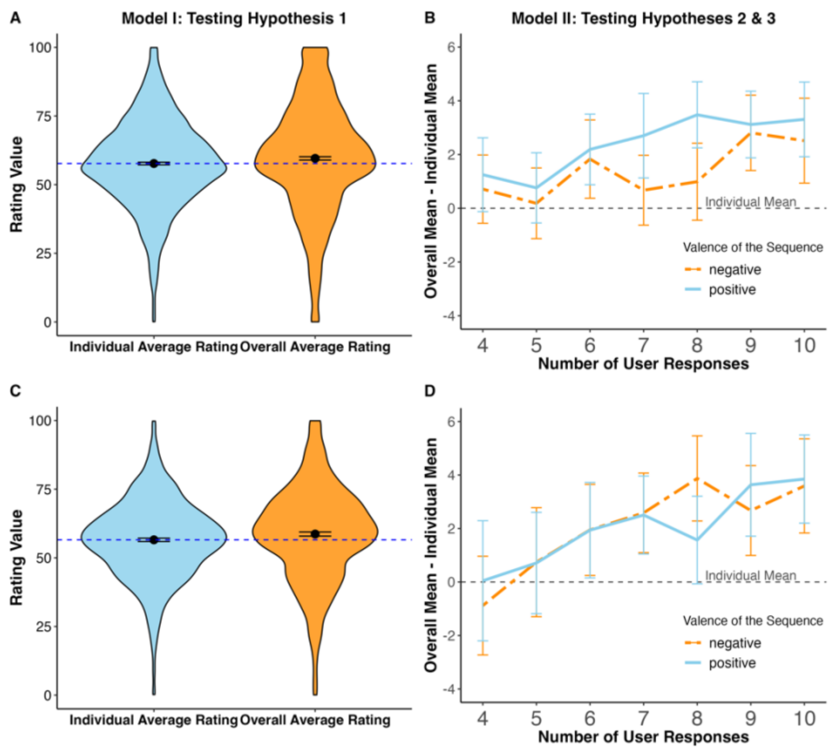

Figure 4. Results of Study 1a (panel A & B) and Study 1b (panel C & D)**.** The violin plot (A & C, Hypothesis 1) shows the comparison between the individual average of a sequence and participants' estimations of the average emotional intensity. The dots represent the means, and the whiskers indicate the 95% confidence intervals. The blue dotted line marks the midpoint of the emotional intensity of the stimuli. For the line plots (B - D, Hypotheses 2 - 3), the x-axis represents the number of responses in the sequence. The y-axis represents the difference between participants' estimation of the estimated average of the response sequence and the individual average emotional intensity of the responses in each sequence. Data are presented as mean value +/- confidence intervals. The grey dotted line represents the mean rating of individual responses.

Overestimation of the Sequence and Willingness to Share News Article. In Study 1b, we tested whether overestimation of the sequence predicted an increased hypothetical willingness to share the news article. To examine this, we used a multilevel model with random intercepts for both story and participant, predicting willingness to share based on the difference score between the estimated emotional instensity average and the individual average ratings of the sequence. The results indicated that greater overestimation was associated with a higher willingness to share (b = 0.17 [0.51, 0.61], SE = 0.048, t (1956.615) = 3.58, p < .001, R^2 = .004). In the SI, we also test how valence, and news article intensity predicts sharing willingness across Studies 1b, 2, and 4.

Discussion

The results of Studies 1a and b suggest that participants tended to overestimate the estimated average emotional intensity expressed in sequences of responses compared to the average of the individual ratings and that this tendency was stronger when increasing the length of the sequence. We did not find significant differences between positive and negative sequences. The emotional intensity of the reaction to the news articles impacted the degree of overestimation, with more intense news articles leading to greater overestimation. Additionally, Study 1b demonstrated that this overestimation has an impact on the emotional reaction intensity to the news article. Moreover, these amplified perceptions had consequences for behavioral intentions: participants reported a higher willingness to share news articles associated with response sequences where overestimation occurred. Next, we investigate the underlying mechanism behind this phenomenon in the following study.

Study 2: Testing memory as the mechanism for overestimation of emotions in sequences of responses to news articles.

Our next goal was to explore potential mechanisms for this overestimation. In Study 2, we examined whether memory for emotionally intense responses contributed to the observed overestimation effect. Based on previous research on emotional memory (Kensinger, 2020; Kensinger & Schacter, 2008), we hypothesized that participants might not recall all items in a response sequence when making their average judgment but would instead be more likely to remember the more emotional responses. This selective memory could then lead participants to perceive the average emotional intensity of the sequence as higher than it actually was. To test this, we examined whether participants remembered emotionally intense responses better than less emotional ones. If memory is indeed biased towards intense responses, it could influence how people aggregate emotional information, potentially accounting for the overestimation effects observed in Studies 1a and 1b.

Method

Stimuli

News Article and Text Responses*.* We used the same 40 news articles and corresponding 15 responses as in Study 1b.

Participants

The power analysis using a pilot sample of 50 participants indicated that we needed to recruit 200 native English speakers from the United Kingdom and the United States to achieve 80% power (see SI). Participants were recruited from Prolific. In line with our pre-registered criteria, we excluded 2 participants who failed more than 3 attention checks or had missing data. Additionally, we excluded 7 participants in line with our pre-registered exclusion criteria, which was identical to Study 1a and 1b. Our final sample consisted of 191 participants (95 females, 94 males, age: M = 39.90, SD = 13.68, range = 18-70), who received compensation of $6.67 for their participation in the 50-minute task.

Procedure

Sequential responses (1st block) and memory task (2nd block). The task included two blocks and 30 trials in total. The first block aimed to replicate the results from studies 1a and 1b (see SI for the results), while the second block aimed to demonstrate the role of memory in this overestimation. The first block consisted of 15 trials and followed an identical procedure to Study 1a, with two exceptions: participants were also asked about their willingness to share the news article at the end of each trial (similar to Study 1b), and all sequences consisted of eight responses (instead of varying numbers of responses). We chose to use a fixed sequence length in order to keep the difficulty level of the subsequent memory task constant.

In the second block, participants completed 15 trials designed to assess their memory for previously seen responses. Each trial began with participants viewing one of the fictional news articles from the first block for 10 seconds. Next, they were shown all 15 text responses related to that article used in Study 1b. Of these 15 responses, eight had been presented in the first block, while the remaining seven were new and unfamiliar to the participants. Participants' task was to correctly identify the 8 responses they had seen before from the 15 options. The presentation order of the original and new responses was randomized for each trial. This setup allowed us to measure how the emotional intensity of each response—previously rated by participants in the first block—predicted its likelihood of being accurately recognized in the second block. To ensure the validity of the memory test, we made two preregistered modifications to the original design, as detailed in the Supplementary Information (SI), when it became clear that participants' overall memory performance was high. After completing the second block, participants filled out a brief survey (see SI).

We conducted three iterations of this study to refine the memory task and effectively test its role in overestimation. In the first version, participants saw the responses again in the memory block and had to choose between one previously seen response and one new response, resulting in a ceiling effect (92% accuracy) that made it difficult to detect an association between emotional intensity and memory. In the second version, we increased difficulty by removing the re-presentation of responses and expanding the response choices from two to three, but this reduced statistical power as participants made only one recognition decision per trial. The final version improved power by requiring participants to identify eight previously seen responses from a set of 15, increasing the number of ratings per trial and allowing for a more robust test of the relationship between emotional intensity and memory. The first block, replicating the overestimation effect, remained unchanged across all versions. Full details of these refinements are in the SI.

Results

Emotional Intensity and Memory. Analyzing the results of the second block, we then examined whether the rated emotional intensity of each text response (provided during the first block) predicted the probability of remembering correctly. We argued that participants are more likely to remember responses with emotional intensity that stands out from the sequence. To evaluate this, we needed to consider both how much an item stood out within its sequence and how it compared to a participant's typical ratings. We calculate this 'relative intensity score' by first subtracting the sequence's average intensity from each text response's rating, reflecting the item's intensity relative to the average for that trial. Next, we standardized these scores by converting them to z-scores, indicating how many standard deviations each score deviated from the participant's mean rating. This approach provided a measure of relative intensity that accounted for both the sequence context and individual rating tendencies: *(rating - sequence mean) / standard deviation*. We conducted a generalized linear mixed model using this standardized score as the predictor and whether participants correctly identified this response in the memory test as (dummy-coded) dependent variable. We also included a covariate representing the length of the text response (in number of characters) to control for any influence of this factor on the likelihood of remembering a given text response and added a by-individual random intercept and a random intercept representing each news article in the sequence. Results showed that the relative intensity of the response positively predicted memory accuracy (b = 0.040, z (2.69) = 3.84, p < .001, R^2= .05, 95% Confidence Intervals = [0.02, 0.06).

To further analyze whether overestimation was more pronounced in sequences where the tendency to remember more emotionally intense items was stronger, we initially computed two difference scores for each sequence. First, we calculated the difference between the estimated average emotional intensity rating and the average individual emotional intensity rating of the sequence, serving as a measure of sequence overestimation. Second, we determined the average emotional intensity of the items that participants recalled compared to the average individual ratings of the sequence. A higher score indicated a stronger tendency among participants to remember more intense items from the sequence. Using a linear mixed model, we predicted whether the overestimation of a sequence would be predicted by the extent to which a participant remembered more intense items in each trial, with the article intensity considered as a covariate. We also included random intercepts for participants and stories. However, we found no significant association between the overestimation in the evaluation of emotions in a response sequences and the propensity to remember more intense items in the sequence (b = 0.05, t (2788.04) = 1.37, p = .168, R^2= .19, 95% Confidence Intervals = [-0.02, 0.13]). We also investigated other memory-related concepts, such as how item position in the sequence influences the estimation of the sequence, and the relevant analyses can be found in the Supplementary Information.

Discussion

This study provides initial evidence for a potential mechanism behind overestimation: participants were more likely to remember responses with higher emotional intensity. However, we did not find that participants prone to overestimation were also the ones who remembered more emotionally intense responses. This indicates that while memory may play a role in overestimation, the study does not establish a direct connection between the two. To further investigate whether the better memory of emotionally intense responses can explain overestimation, we will next examine whether overestimation persists in sequences where emotional intensity is the only distinguishing factor between items. If overestimation occurs under these conditions, it would suggest that emotional intensity itself is driving the overestimation effect.

Study 3: Overestimation in the evaluation of sequences of words

The aim of this study was to examine whether emotional intensity alone is sufficient to drive overestimation or whether additional contextual factors in previous studies contributed to the effect. In Studies 1a and 1b, participants evaluated full-text responses to news articles, which contained multiple sources of information, such as sentence structure, response length, and punctuation—elements that may enhance the salience or diagnosticity of emotional expressions. Additionally, the news article itself set a context for interpreting the responses, which we found did influence overestimation. To better isolate the role of emotional salience, we presented participants with sequences of individual emotional words, rather than full responses. While we acknowledge that individual words are not devoid of meaning or contextual associations, they lack the structured narrative context of full responses and do not describe specific events to which emotionality could be attributed. Unlike full-text responses, where informativeness is shaped by narrative coherence, words presented in isolation rely primarily on intrinsic emotional valence to drive salience. By testing whether overestimation still occurs under these conditions, Study 3 serves as a conceptual replication that removes contextual framing and instead presents a sequence of unrelated emotional stimuli. If overestimation persists in this highly controlled setting, it suggests that the effect is not solely a function of the way emotions are conveyed in real-world conversations but rather a more general cognitive tendency when aggregating emotional information.

Methods

Stimuli

Words. We obtained the words for Study 3 from the Evaluative Lexicon 2.0 (EL 2.0), developed by Rocklage, Rucker, and Nordgren (2018), which contains a large set of 1,541 evaluative words that have been rated by approximately 600 participants on their implied valence (0: Very negative, 9: Very positive) and emotionality, i.e., the extent to which they imply an emotional, feelings-based reaction (0: Not at all emotional, 9: Very emotional). Thus, each word has an aggregate valence rating, and an aggregate emotionality rating based of approximately 30 ratings. The words span from very negative ("worst" = .76 out of 9.00) to very positive ("flawless" = 8.24 out of 9.00) on valence, and from relatively low emotion ("healthy" = 2.38 out of 9.00) to high emotion ("enjoyable" = 6.58 out of 9.00) on emotionality. These emotionality ratings were used to calculate the objective emotionality of a sequence in our analysis. To avoid repetitions, we only included word stems in our stimulus set, omitting all word derivatives. For example, both "happily" and "happy" existed in the EL 2.0, but we chose to retain only the word stem "happy" in the final stimulus pool. Additionally, we only included words that were in both the EL 2.0 and ANEW dictionary by Bradley and Lang (1999), because the ANEW dictionary provides information about a word's frequency of use in everyday language, which enabled us to distinguish between the effects of a word's frequency and its intensity rating in our analysis. This step narrowed down our selection to 853 words that were presented to participants in the task.

Participants

Based on our power analysis (see SI), we recruited 100 native English speakers from the United Kingdom and the United States via Prolific (66 females, 33 males, 1 other, age: *M* = 41.86, *SD* = 15.39, range = 18-74). Specifically, in this study they were compensated $2.30 for their participation in the 20-minute task. The exclusion criteria applied were failing three or more attention checks or having incomplete information in any of the 50 trials or in the end survey questionnaire. All participants were retained in the final sample.

Procedure

Sequential responses task. The task consisted of 3 practice trials, which were excluded from the analysis, followed by 50 actual trials. In each trial, participants were presented with a sequence of 1–15 words, with the number of words per trial randomly determined at the start of each trial. Each word was displayed individually in the center of the screen for 1000ms, followed by a fixation cross lasting 400–600ms before the next word appeared. Unlike the previous studies, we decided to include sequences of fewer than 4 items to demonstrate that participants do not overestimate short sequences. The corpus of words consisted of adjectives and nouns that varied in emotional intensity, but the valence within each sequence was not mixed. Specifically, in the negative condition, all words in the sequence ranged from more neutral to negative, while in the positive condition, they ranged from more neutral to positive. More neutral words were still classified as either positive or negative but had low, non-zero emotional intensity ratings. For example, the negative-valenced word "*discretion*" had a low emotional intensity rating of 0.02, while the lowest-intensity positive-valenced word, "*homey*," had a rating of 0.1. Across the stimuli, negative words ranged from 0.02 to 3.72 in emotional intensity, while positive words ranged from 0.1 to 4.07. This demonstrates that some words in the sequence were only slightly emotional, while others were more extreme in intensity. The words in the sequence were randomly selected from a list of 853 stimuli and displayed on a white background in the center of the screen for 1000ms. For full details on the specific words used and their emotional intensity ratings, see Supplementary Information (SI). Participants could see each of the words only once in all trials. After each word, participants saw a fixation cross for a duration between 400-600ms (randomly determined on each trial). After each sequence, participants were asked to "*estimate the average emotion expressed by the sequence of words*" on a scale from 0 (Not at all Emotional) to 9 (Very Emotional) without any time restriction. After completing the main task, participants filled out a short survey including questions about demographics, personality and social anxiety (for full detail, see SI).

Results

Evaluating Overestimation of Response Sequences. We first examined whether participants' estimation of the overall average emotional intensity of the sequence was higher than the average of the pre-computed evaluations of the individual words in the sequence (H1). We conducted a mixed model repeated measures analysis, comparing the average individual emotional intensity of the words in each sequence with participants' estimation of the overall average emotional intensity of the sequence. As the frequency of a word occurring in natural language could potentially influence its salience, we added the frequency score as a covariate. These frequency scores were sourced from the ANEW (Affective Norms for English Words) dictionary (Bradley & Lang, 1999). Additionally, we incorporated a random intercept for each participant in the model. In support of the initial hypothesis (H1), the estimated average emotional intensity of the sequence was 0.75 scale points higher than the average individual emotional intensity of the words in each sequence, as measured on a 10-point scale (b = 0.76, SE = 0.032, t (9893.93) = 23.53, $p < .001$, $R^2 = .05$, 95% Confidence Intervals = [0.69, 0.82]). We then tested our second hypothesis that overestimation of average emotion evaluation would be greater for longer sequences, and our third hypothesis, which suggested that negative sequences would exhibit higher levels of overestimation. We used the same multilevel model as in Studies 1 and 2, predicting the difference between the estimated and actual average emotional intensity based on sequence length and valence (positive vs. negative). The model included a random intercept for each participant. Our results showed that overestimation of the average emotion of the sequence was greater when the sequence of words was longer (b = 0.08, SE = 0.0085, t (4903.56) = 9.55, $p < .001$, $R^2 = .06$, 95% Confidence Intervals = [0.06, 0.10], see Figure 5). Importantly when a sequence contained only one word (represented by the intercept of the model), the difference between the pre-computed intensity score and the participant's estimation was not significantly different from zero (b = -0.19, SE = 0.14, t (188.06) = -1.37, $p = .17$, $R^2 = .06$, 95% Confidence Intervals = [-0.47, 0.08]). This absence of overestimation in emotional intensity ratings for single-word sequences supports the alignment between the pre-evaluated emotional intensity scores and our participants' overall ratings, indicating that the overestimation observed in longer sequences is a consequence of sequential presentation rather than discrepancies between pre-evaluated scores and participant ratings. Unlike in the previous studies, our test of the third hypothesis, using the same model, showed that overestimation of emotional intensity in sequences was significantly greater for negative sequences than for positive sequences (b = 0.25, SE = 0.11, t (4902.43) = 2.26, p = .023, $R^2 = .06$, 95% Confidence Intervals = [0.03, 0.47]).

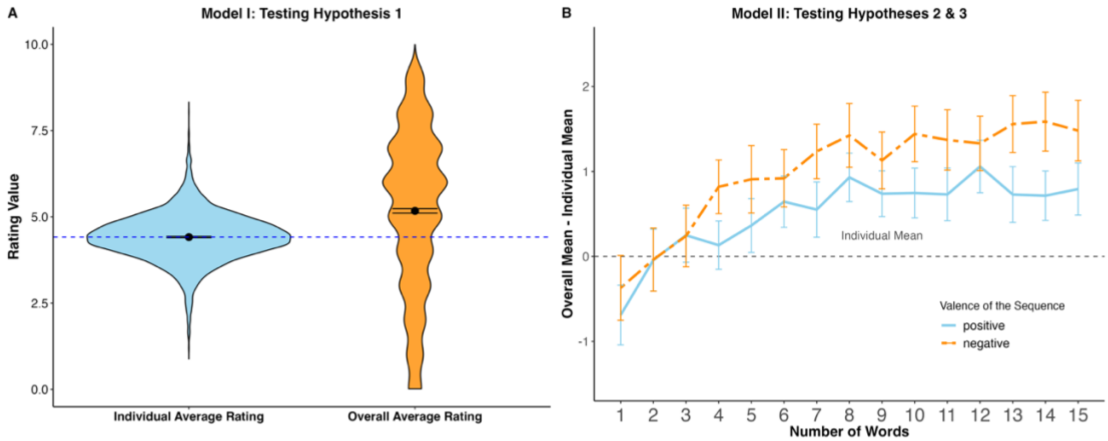

Figure 5. Results of Hypotheses 1 - 3 in Study 3. The violin plot (A, Hypothesis 1) shows the comparison between the individual average of a sequence and participants' estimations of the overall average. The dots represent the means, and the whiskers indicate the 95% confidence intervals. The blue dotted line marks the midpoint of the emotional intensity of the words. For the line plots (B, Hypotheses 2 - 3), the x-axis represents the number of words in the sequence. The y-axis represents the difference between participants' estimation of the overall average of the response sequence and the individual average emotional intensity of the responses in each sequence. Data are presented as mean value +/- confidence intervals. The grey dotted line represents the mean rating of individual responses.

Discussion

Finding overestimation in sequences that reduce text to its minimal form—using random words unrelated to any specific topic rather than comments on a particular news story—suggests that the mechanism behind overestimation is the presence of a sequence with varying salience. Furthermore, the study shows that participants are generally accurate when evaluating smaller sequences. Unlike previous studies, we found that overestimation was greater for negative sequences compared to positive ones. After establishing and exploring the overestimation effect in emotional words and text sequences, we next focus on its implications and consequences.

Study 4: Effects of overestimation on the perception of representativeness

The goal of Study 4 was to test the association between the overestimation of emotions in sequences and participants' perceptions of the representativeness of more emotionally intense responses. Previous research has shown that perceptions of representative opinions and views are often exaggerated on social media compared to offline contexts (for an overview, see Robertson et al., 2024). Here, we aim to test a novel additional mechanism underlying this phenomenon. The salience of emotionally intense responses might influence not only overestimation, as shown in Study 2, but also participants' representation of how others would express themselves in response to the news articles. Specifically, if individuals base their judgment of the general public's opinion on their estimation of the average response they have observed, it might lead them to believe that more emotionally intense responses are more representative of others' opinions. Further, this belief may in turn lead individuals to be more willing to share emotionally intense news articles.

Methods

Stimuli

News Article and Text Responses*.* We used the same 40 news articles and corresponding 15 responses as in Study 1b and Study 2.

Participants

Previous studies using facial stimuli have shown that the overestimation of the sequence is weaker when the number of stimuli is consistent between sequences (Goldenberg et al., 2022). Because Study 2, the only other study using a fixed sequence length, was run in parallel to Study 4, we lacked information to estimate the effect size for Study 4. Instead, we calculated the sample size reported in the pre-registration based on data from Study 1a, focusing on sequences with 8 text responses for a comparable analysis. To achieve 80% power, assuming a similar effect size of Study 4 as Study 1a, we determined that 50 participants completing 15 trials each would be necessary. We recruited 150 native English speakers from the United Kingdom and the United States through Prolific, more than suggested by the power analysis based on Study 1a (see SI). We removed 4 participants in line with our pre-registered exclusion criteria. The final sample comprised 146 participants (75 females, 71 males, age: M = 41.95, SD = 14.90, range = 20 -77). The study took approximately 50 minutes and participants received compensation of $5.75.

Procedure

Sequential responses (1st block) and representativeness assessment (2nd block). This task, similar in structure to the memory task in Study 2, was divided into two blocks, each consisting of 15 trials, with 2 practice trials excluded from analysis. The first block was identical to Study 2, designed to test overestimation of emotional intensity in response sequences (see SI for results). In the second block, participants were presented with the same fictional news articles and response sequences they had previously seen in Block 1. However, instead of rating the estimated average emotional intensity of the sequence, they were now asked to identify the response they believed was most representative of the public's emotional reaction. To do so, they were presented with three text responses to the article, displayed one above the other on the screen. These responses were randomly selected from the sequence participants had seen in the first block, and their order of presentation was also randomized. Participants were instructed to select "*which response you think is most representative of the emotion that others would express in response to this article*" by clicking on it. We framed this question in terms of “representativeness” instead of “normativity” to ensure participants could intuitively understand and respond to the task.

Results

Emotional Intensity and Representativeness. For the analysis of the second block, we calculated the likelihood of a response being selected as most representative of public opinion based on its individual emotional intensity rating. Our analysis used a generalized linear mixed model predicting whether or not the response was selected (dummy-coded) by participants, with the independent variable being its individual emotional intensity rating. We found that participants perceived emotionally intense responses as more representative of how other people would respond to the news article (b = 1.012, SE = 0.017, z = 12.79, p < .001, R^2 = .035, 95% Confidence Intervals = [0.010, 0.014]). This means, according to this model a response with an emotional intensity rating of 90 is approximately 1.61 times more likely to be selected as representative compared to a response with an intensity of 50.

To investigate whether overestimation was more pronounced in sequences where participants tended to perceive the more emotionally intense responses as representative, we initially computed the difference between the estimated average emotional intensity rating and the individual emotional intensity average of the response sequence, serving as a measure of sequence overestimation. Using a linear mixed model, we examined whether the overestimation of a sequence could be predicted by the emotional intensity of the item participants selected as most representative of this sequence, with the article intensity considered as a covariate. Additionally, random intercepts for participants and stories were included. Our analysis revealed greater overestimation in the evaluation of emotions in response sequences when the emotional intensity of the most representative item in the sequence was higher (b = 0.03, t (2788.04) = 1.37, p = .006, R^2 = .003, 95% Confidence Intervals = [0.01, 0.05]).

Representativeness of Response as a Predictor for Willingness to Share News Article. We investigated how the overestimation of the sequence's emotional intensity and the emotional intensity of the most representative response influenced participants' hypothetical willingness to share the news article. Using a linear mixed model with random intercepts for participants and news stories, we predicted participants' willingness to share based on the difference score between the estimated overall average and the individual average ratings of the sequence (as a measure of overestimation) and the emotional intensity of the most representative response within the sequence. Our results showed that both factors significantly and independently contributed to participants' sharing tendencies. Specifically, the degree of overestimation was a significant predictor (b = 0.20, SE = 0.023, t (2096.09) = 8.12, p < .001, η^2 = .01, R^2 = .03, 95% Confidence Intervals = [0.13, 0.27]), as was the emotional intensity of the most representative response (b = 0.16, SE = 0.73, t (2141.50) = 3.88, p < .001, η^2 = .03, R^2 = .03, 95% Confidence Intervals = [0.12, 0.20]).

Discussion

The finding of this study reveals important consequences of how people sample and aggregate emotions on social media. Specifically, the way individuals sample and interpret information affects not only their estimation of the average response but also their perceptions of representativeness, such as which responses are seen as most reflective of others' emotional reactions. This process can shape behavioral intentions, as participants reported being more willing to share content they perceived as eliciting stronger emotional reactions from others. These intentions may, in turn, perpetuate the perception of heightened emotional intensity. In our final study, we aim to determine whether this overestimation effect extends beyond individual sequences to influence perceptions of the entire newsfeed.

Study 5: Evaluating the Emotional Intensity of Entire Newsfeeds

The studies reported so far have shown participants only one response at a time and presented news stories separately in sequence. Although these design decisions helped to maximize experimental control, they also reduced generalizability. Specifically, when users engage with social media outside the laboratory they are often exposed to several overlapping streams of information. Study 5 was designed to assess whether participants would overestimate the emotional intensity of an entire newsfeed rather than being confined to individual news articles and comments in isolation. Additionally, this study provided an opportunity to address an important ecological validity concern: whether overestimation still occurs at an aggregate level when a sequence includes mixed-valence content in terms of the news articles, while the comments within each article remain consistent in valence. This design mirrors a key characteristic of real-world social media feeds, where the emotional tone of news articles often varies, even if comments within each thread align with the article's valence. Thus, the main goals of Study 5 were to evaluate participants' emotional assessments of (1) text sequences and (2) an entire newsfeed composed of mixed-valence news articles and same-valence comments, mirroring real-world social media.

Methods

Stimuli

News Article and Text Responses. We used the same 40 news articles and corresponding 15 responses as in Study 1b, Study 2 and Study 4.

Participants

We recruited 200 native English speakers from the United Kingdom and the United States through Prolific. Since this study had a different design and analysis approach, we could not use power analysis to estimate the required sample size. Instead, we chose to recruit a number of participants equivalent to that in the study with the smallest effect sizes (Study 2) to ensure sufficient power for detecting similar effects. Following our pre-registered exclusion criteria, which were identical to all previous studies, we removed 3 participants. The final sample consisted of 197 participants (97 females, 97 males, and 3 others, age: M = 43.77, SD = 13.33, range = 19 -75). The study took approximately 40 minutes and participants received compensation of $5.50.

Procedure

In each trial of this study, participants saw a layout resembling a social media newsfeed beginning with a fictional news article presented in the style of an X-post (formerly Twitter). Beneath this post was a scale for participants to rate their emotional reaction in response to the news article, ranging from 0 (neutral) to 100 (very emotional) as in the previous studies. Text responses were displayed as comments below the scale. Unlike previous studies, we refrained from asking participants to rate each response individually. This approach aimed to provide a more realistic assessment of how users might engage with content on their actual social media newsfeed, as users may not consistently make individual judgments for all items without being prompted to do so. Consequently, we were unable to calculate individual averages. Each news article post included between 1 and 15 responses. Each news article post included social media features such as like and share buttons, a displayed comment count, and an author's name.

These features were functional, allowing participants to click on share or like, but we did not instruct them to do so and thus did not analyze their interaction with these features. The authors' names were randomly chosen from a list of 25 male and 25 female names, ensuring no repetition within a single newsfeed. Beside each author's name was a user icon, represented by a silhouette in one of eight randomized colors. The posts contained between 1 and 15 responses. We expanded the range of displayed responses from the 4-10 used in previous studies to maximize variance, as sequence length was the primary independent variable in our study. Below the comment section was another slider, also ranging from 0 (neutral) to 100 (very emotional), and participants were asked to estimate the average emotionality of the text responses by asking "*Please rate the average emotions expressed in the comments to the news story in the comment section*". To ensure that participants did not simply scroll to the end of the page or view content in a mixed order, they were only able to proceed to the next set after rating both the news article and the comments. Only after participants had rated both the news article and the comments, they advanced to the next set, consisting of a new news article post, a news article rating slider, a comment section, and a comment section rating slider. Each newsfeed contained both positive and negative news stories, with corresponding responses in their comment sections. The responses within each comment section matched the valence of the associated news article as in the previous studies. Each trial involved rating a total of six such sets, except in the practice trial, where only three sets were rated. After viewing the entire newsfeed, participants were asked to rate the average emotional intensity of the newsfeed on a slider ranging from 0 (neutral) to 100 (very emotional) with the prompt: "*Please rate the total average emotionality of this entire news feed*". All sliders started at 0 to anchor the ratings towards the most conservative estimate of emotionality. Participants completed 6 trials in total, as well as one practice trial. Upon completing the primary task, participants were directed to an identical survey as in the previous studies.

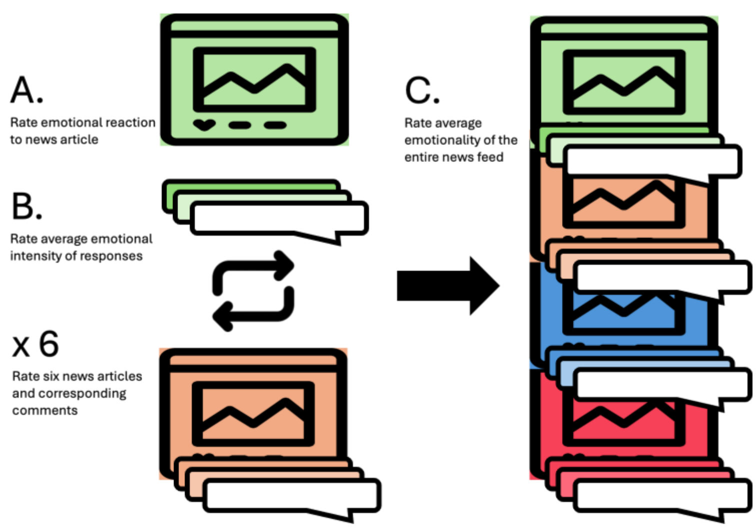

Figure 6. Design for Study 5. Participants rated their emotional reaction to the news article (A) and then the average emotional intensity expressed by the corresponding response sequence (B). This process was repeated for six pairs of news articles and their corresponding responses, presented consecutively in a simulated newsfeed. After completing these ratings for all pairs, participants were asked to rate the average emotional intensity of the entire newsfeed they had just viewed.

Results

Evaluating Estimation of the Newsfeed and Responses. We examined whether individuals tended to overestimate the emotional intensity of entire newsfeeds rather than focusing solely on responses to individual news stories (H1), as in previous studies. To assess this, we compared the estimated average emotional intensity of the entire newsfeed, the average of the ratings of all presented news articles, and the average of the estimated intensities of all presented response sequences. Participants provided a direct rating of the overall emotional intensity of the entire newsfeed using a 0–100 scale after viewing the full set of articles and responses. The average news article intensity was calculated based on participants' individual ratings of each news article, while the average response sequence intensity was based on participants' overall emotional intensity ratings of each comment section. Overestimation at the newsfeed level occurred when participants rated the emotional intensity of the entire newsfeed as higher than both the average ratings of individual news articles and the average of estimated intensities of response sequences. Testing this comparison, we used a mixed-model analysis of repeated measures, incorporating random intercepts for participants. Our findings revealed that participants rated the entire newsfeed as more emotionally intense compared to the individual average of the news articles (b = -7.43, SE = 0.48, t (3346.59) = -15.20, p < .001, R^2 = .03, 95% Confidence Intervals = [-8.38, -6.47], see Figure 7A) and the average estimated intensity of the individual response sequences (b = -3.97, SE = 0.48, t (3346.56) = -8.15, p < .001, R^2 = .03, 95% Confidence Intervals = [-0.33 – 0.01]). To test if the sequence length influences participants' estimations of average emotional intensity (H2), and whether this increase differs between positive and negative news articles (H3), we used a single model. In this model, we predicted participants' estimations of the average emotion expressed in responses to news articles in comments, using sequence length, and its interaction with valence as predictors. We included the rating of the intensity of the emotional reaction to the news article as covariates and a random intercept for each participant and news story. Our results indicated that participants perceived longer sequences as more intense (H2: b = 1.23, SE = 0.061, t (6944.62) = 19.92, p < .001, R^2 = .29, 95% Confidence Intervals = [1.11 – 1.35], see Figure 7B). However, the increase in estimated intensity did not significantly differ between positive and negative news stories (H3: *b* = -0.16, *SE* = 0.032, *t* (9893.93) = -1.81, *p* = .07, R^2 = .29, 95% Confidence Intervals = [-0.33 – 0.01]). Additionally, in the supplementary materials, we provide a detailed analysis of how positive and negative news articles, along with their corresponding responses, contribute to overall newsfeed intensity ratings.

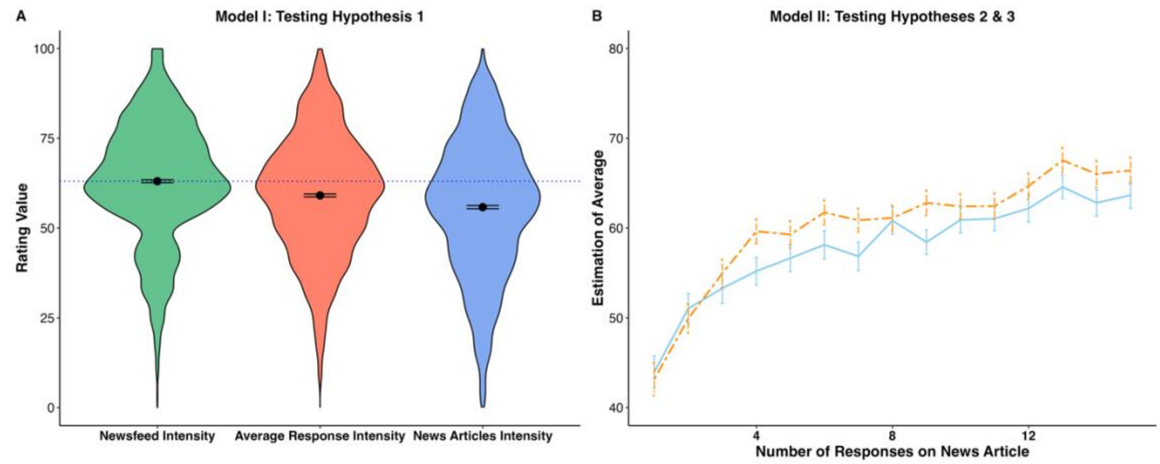

Figure 7. Results of Hypotheses 1 - 3 in Study 5*.* The violin plot (A, Hypothesis 1) shows the comparison between the overall news feed rating, the average intensity of the reactions to the news articles, and the average of the responses to the news articles. The dots represent the means, and the whiskers indicate the 95% confidence intervals. The blue dotted line marks the emotional intensity rating of the overall news feed. For the line plots (B, Hypotheses 2 - 3), The x-axis represents the number of responses in the comment section under a news article. The y-axis represents the participants' estimation of the average sequence. Data are presented as mean value +/- confidence intervals.

Discussion

When evaluating the emotionality of the newsfeed as a whole, participants again tended to overestimate its intensity compared to individual news stories and their responses. This suggests that social media users not only overestimate how others feel and express themselves about specific topics but also perceive the overall sentiment across the platform as more extreme than it actually is, further distorting their perception of reality. Additionally, this study demonstrates that emotional intensity can be overestimated even when the content includes mixed-valence news articles. This overestimation may license users to share or express themselves in more extreme ways if they believe such behavior aligns with the perceived norm on the platform.

General Discussion

Across six studies, we consistently found that people overestimated the emotional intensity of sequences (vs. individual texts) (H1). This overestimation was stronger for longer sequences (H2), and for more intense news articles. We did not find systematic differences in overestimation between positive and negative sequences of responses to news articles (H3). We also found that overestimation of the sequence predicted participants' willingness to share the news article. To better understand the mechanisms underlying this effect, we first tested the hypothesis that emotionally intense responses are more likely to be remembered. While our findings confirmed that participants were more likely to remember emotionally intense responses, we did not find a direct link between this tendency and overestimation. Additionally, to demonstrate that the overestimation is driven by the salience of emotionality, we replicated the overestimation of sequence averages using sequences composed solely of individual words, showing that overestimation results from emotional content rather than other contextual cues. Analyzing the consequences of overestimating the emotionality of text sequences as encountered on social media, we showed that overestimating sequence averages was also associated with judging that the more intense responses better represent how the public would react to them and that participants also overestimate the average emotional intensity of entire newsfeeds.

The effect size of overestimation in response sequences ranged from 1.86% to 2.12% per news article. While small in isolation, this effect gains practical significance given the vast amount of content users encounter daily of around 100 posts (Piccardi et al., 2024). Additionally, as demonstrated in Hypothesis 2, overestimation increases with the amount of content consumed sequentially. In Study 5, participants evaluating a longer feed perceived the feed's emotional intensity as 7.43% higher than the average intensity of the news articles and 8.15% higher than the average intensity of the corresponding comment sections. Notably, this overestimation emerged after viewing just six news articles, each accompanied by 4–10 responses. Even small overestimations, when applied across this volume of content, can compound to significantly amplify perceived emotional intensity over time

One important question is whether these overestimations truly constitute a bias or not. One way to define bias as a systematic deviation from accuracy based on the instructions given (Lerner et al., 2015). Given that participants were explicitly asked to estimate the average emotional intensity of a sequence, their judgments were systematically skewed toward perceiving sequences as more intense than their own individual response ratings would suggest. In that sense, we observed a bias. Additionally, we observed that more recent emotional content disproportionately influenced the estimated average emotional intensity rating, aligning with both the emotional immediacy bias (Van Boven et al., 2009) and recency bias (Tversky & Kahneman, 1973). Others sometimes define bias in terms of deviations from rationality or even maladaptiveness (Gigerenzer, 2018). From an ecological rationality standpoint, what appears to be a bias in laboratory tasks may actually represent adaptive responses to patterns we typically encounter. From this perspective, whether overestimation is a bias is more nuanced. On the one hand, although our observed overestimations deviate from strict accuracy, they can still be advantageous (for an overview, see Lerner et al., 2015). For instance, if a news article garners numerous reactions, that volume may legitimately indicate heightened salience or public interest (Effron & Brady, 2025), making overestimation adaptive to efficiently identify socially relevant information. On the other hand, social media platforms often exploit these mechanisms in ways that may not map onto true importance—viral trends, echo chambers, and algorithmic amplifications can inflate engagement without corresponding to actual significance. Under these conditions, repeated overestimation can be maladaptive by distorting perceptions of reality, increasing stress, and fueling polarization. Consequently, although overestimation constitutes a real bias by the strict definition of "systematic inaccuracy," it is not inherently maladaptive. Instead, it can evolve from a helpful heuristic into a harmful distortion, depending on how well the environment's cues align with genuine emotional salience.

Contrary to our second hypothesis, we did not find systematic differences in overestimation between positive and negative sequences in responses to news articles (H3). Only in Study 3, which used individual words rather than full-text responses, was overestimation stronger for negative sequences than for positive ones. Based on negativity bias, we initially hypothesized that negative content would be more salient and memorable (Baumeister et al., 2001; Rozin & Royzman, 2001). There is prior evidence from research on overestimation of emotionality in faces suggesting such valence differences, but when observed, these effects have generally been weak and inconsistent (Goldenberg et al., 2021, 2022). To increase the likelihood of detecting such an effect, we pooled data from Studies 1a, 1b, 2, and 4, but still did not observe a significant difference between positive and negative sequences (see SI). One possible explanation is that text stimuli introduce additional complexity compared to facial expressions, as factors such as length, narrative structure, and personal relevance may interact with valence in ways that make small effects more difficult to detect. Although we did not observe a valence difference when analyzing responses to the same article, our analysis of whole-feed estimations in Study 5 suggested that negative comments exerted the strongest influence on overall newsfeed intensity ratings (see SI). This finding raises the possibility that detecting valence effects in text-based stimuli may require exposure to a larger volume of text within each trial, a question that should be explored further in future research.

These findings contribute to the growing literature on how emotions are amplified on social media (Goldenberg & Willer, 2023). In a recent study, Brady et al. (2023) found that individuals who view social media content tend to perceive emotions in texts as stronger compared to the estimation of those who produced the texts. Here, we uncover an additional process of overestimating emotional content driven by how people assess sequences rather than individual pieces of content. These two forms of overestimation may be connected, as individuals may rely on their overall experience of the overall average emotional intensity of social media content to evaluate individual pieces and vice versa. If users indeed utilize their perception of the average to inform their own judgments, it could lead them to perceive responses as more emotional than intended. Given that much of the text that people produce online is presented and evaluated in sequences, understanding overestimation driven by the aggregation of information provides another significant layer of insight into how individuals interpret others' emotions online.

Our findings could have significant implications for users' behavior online by pushing users to express stronger emotions. Overestimation in the evaluation of emotions in text sequences may impact the content users share and produce. Because people tend to overestimate the emotional intensity of responses to news articles, as well as the entirety of the content they encounter in a newsfeed, their perception of what types and intensities of emotions are more prevalent (descriptive norm) and acceptable (injunctive norm) could be influenced. Our studies already demonstrated that the estimation of the average does increase willingness to share content. Additionally, Study 4 provided initial support for this claim by demonstrating an association between the overestimation of the response sequence emotionality and participants' perceptions of what was most representative in the sequence. This misperception not only affects users' perceptions but also encourages them to share and produce more emotionally intense content themselves, which should be examined in future research.

Overestimation of emotions expressed in sequences might also have well-being consequences. The content we consume has a direct impact on our emotions and, consequently, our overall well-being (Kelly & Sharot, 2023). Heightened perception of negativity, either of overall responses or their corresponding events, can lead to increased worry and anxiety about those events (Kort-Butler & Hartshorn, 2011; Mello et al., 2022; Näsi et al., 2015). Developing an exaggerated perception of the average positivity can be detrimental too. When individuals perceive positive content as more positive than warranted by the average positive content they encounter, they may engage in unrealistic comparisons with their own lives resulting in a feeling of inadequacy or dissatisfaction (Collins, 1996). These ideas should be further examined in future research.

Limitations and Future Directions.

Despite the important implications of our findings and our efforts to address alternative explanations, there are limitations to the studies that may hinder the generalizability of our findings (summarized in the Table of Limitations). The first notable limitation pertains to the external validity of our findings. In our study, participants were provided with content without any background information on the creators, and the events in the news articles were entirely fictional, giving them no opportunity to seek further details. We chose this design to reduce noise and reveal the desired findings. However, in real social media settings, users often have prior knowledge about the content creators and may not treat all responses equally; they may tend to give more weight to content from individuals they know, trust, or recognize, such as public figures (J. P. Schöne et al., 2023). This selective attention could influence their perceptions of events and other responses, potentially leading them to not overestimate the averages as much as suggested in our study. Future research should examine the impact of prior knowledge and the identity of the content producer on content perception and public opinion.

A second limitation of the current finding is that in our study, the number of responses was randomized, whereas on social media, the number of responses to a news article typically signals the general public's interest and perhaps their emotionality in response to the article. Users on social media, might infer that a news article with many responses has greater emotional intensity, as a large number of comments suggests widespread desire to express their opinions and emotions. While sequence length might have informed people's estimation of the average emotion expressed in Studies 1a and 1b, and especially in Study 5, where the number of all responses was visible because they were presented together, we still observed overestimation in the estimation of averages in that used words without context (Study 3), where the sequence length does not inherently imply meaning such as public interest. To investigate the influence of numerical information on the evaluation of emotions in the news article further, future research could display the number of comments as a simple count under the content, mimicking how it is presented on social media, to gauge its effect.

A third limitation arises from the random selection of responses to the news articles from the available set. In real-life situations, emotional responses may be more clustered and align more closely with the emotional intensity of the news article itself. For instance, when confronted with an extremely emotional news piece, the likelihood of encountering multiple neutral responses would diminish, with most reactions likely falling within a similar emotional intensity range. This clustering of responses simplifies the task of assessing the average emotional intensity. If individuals do encounter neutral content, they might not include it when calculating the emotional average, considering it less credible or relevant. This difference in the distribution of emotional intensity between our study (normally distributed) and real social media (potentially narrower distribution) could contribute to the higher degree of average response overestimation observed in our study compared to what might be expected in real social media contexts.

Finally, a fourth limitation relates to the ecological validity of the study, as responses to the news articles always matched the valence of the news articles themselves. In real-world contexts, comments often vary not only in emotional intensity but also in, and some may even explicitly disagree with the news article. We chose not to include mixed-valence responses or disagreement within individual sequences to isolate the role of emotional salience by systematically varying emotional intensity while holding valence and agreement constant. Introducing variability in valence or disagreement would have introduced additional complexity to the aggregation process, making it unclear whether opposing-valence comments would be disregarded as outliers (Dannals & Miller, 2017) or weighted more heavily due to their distinctiveness (Haberman & Whitney, 2010). This ambiguity would make it difficult to predict in which circumstance overestimation may occur, which would require having multiple additional studies to build a theoretical framework. Study 5, which featured a news feed with both positive and negative stories, allowed us to provide evidence that overestimation still occurs at the feed level when valence is mixed across different articles. However, it remains an open question for future research how people estimate averages when both valence and agreement vary within a single sequence.

To conclude, our findings suggest that people tend to overestimate the average emotions presented in sequences of words and texts. This insight helps us understand how individuals perceive emotions expressed by others online given that most websites, such as social media, are designed sequentially. Understanding how people assess the average emotions conveyed in sequences, such as texts, could play a crucial role in explaining both how they experience emotions themselves, produce and share content online, and how they perceive the world beyond social media.

Table of Limitations

Limitation	Limitation Details
Lack of Prior Knowledge and Familiarity with the Content and Creator	Participants were exposed to fictional news articles without background information or context about the content creators. In real social media settings, users' prior knowledge and familiarity with content creators can influence their perceptions, potentially reducing the overestimation observed in our study.
Randomized Number of Responses	The number of responses to each news article was randomized, whereas, on social media, the number of responses often signals public interest and emotional intensity. This may influence perceptions differently compared to real social media environments.
Artificial Distribution of Emotional Responses	Responses were randomly selected, which may not reflect real-world scenarios where emotional reactions are often clustered and aligned with the intensity of the news content. This could lead to more accurate estimations of emotionality in real settings compared to our study findings.
Hypothetical Sharing Behavior	The study used an experimental social media platform, allowing us to measure only hypothetical sharing willingness. It is unclear how participants would behave on their actual social media accounts, potentially affecting the ecological validity of the findings.
Fictional Content and Real-World Impact	Since the events in the study were fictional, we could not assess how overestimation of emotions and norms might affect participants' real beliefs and well-being.

References

Bates, D., Maechler, M., & Bolker, B. (2013). *Lme4: Linear mixed-effects models using S4 classes. R package version 0.999999-0. 2012*. http://cran.r-project.org/package=

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, *5*(4), Article 4. https://doi.org/10.1037/1089-2680.5.4.323

Becerra, R., Preece, D., Campitelli, G., & Scott-Pillow, G. (2019). The Assessment of Emotional Reactivity Across Negative and Positive Emotions: Development and Validation of the Perth Emotional Reactivity Scale (PERS). *Assessment*, *26*(5), 867–879. https://doi.org/10.1177/1073191117694455

Bradley, M. M., & Lang, P. J. (1999). *Affective norms for English words (ANEW): Instruction manual and affective ratings*. Technical report C-1, the center for research in psychophysiology ….

Brady, W. J., Crockett, M., & Bavel, J. J. V. (2019). *The MAD model of moral contagion: The role of motivation, attention and design in the spread of moralized content online*. https://doi.org/10.31234/OSF.IO/PZ9G6

Brady, W. J., Jackson, J. C., Lindström, B., & Crockett, M. (2023). *Algorithm-Mediated Social Learning in Online Social Networks*. OSF Preprints. https://doi.org/10.31219/osf.io/yw5ah

Brady, W. J., McLoughlin, K., Doan, T. N., & Crockett, M. J. (2021). How social learning amplifies moral outrage expression in online social networks. *Science Advances*, *7*(33), eabe5641. https://doi.org/10.1126/sciadv.abe5641

Brady, W. J., McLoughlin, K. L., Torres, M. P., Lou, M., Gendron, M., & Crockett, M. J. (2023). Overperception of moral outrage in online social networks inflates beliefs about intergroup hostility. *Nature Human Behavior*. https://osf.io/k5dzr/

Collins, R. L. (1996). For better or worse: The impact of upward social comparison on self-evaluations. *Psychological Bulletin*, *119*(1), 51–69. https://doi.org/10.1037/0033-2909.119.1.51

Dannals, J. E., & Miller, D. T. (2017). General social norm perception in groups with outliers social norm perception in groups with outliers. *Journal of Experimental Psychology: General*, *146*(9), Article 9.

Derks, D., Fischer, A. H., & Bos, A. E. R. (2008). The role of emotion in computer-mediated communication: A review. *Computers in Human Behavior*, *24*(3), Article 3. https://doi.org/10.1016/j.chb.2007.04.004

Effron, D., & Brady, W. J. (2025). *Affective And Cognitive Underpinnings of Moral Condemnation When News of Transgressions Goes Viral*. OSF. https://doi.org/10.31219/osf.io/wyqra_v1

Eimer, M., & Holmes, A. (2007). Event-related brain potential correlates of emotional face processing. *Neuropsychologia*, *45*(1), Article 1. https://doi.org/10.1016/j.neuropsychologia.2006.04.022

Fazio, R. H., Pietri, E. S., Rocklage, M. D., & Shook, N. J. (2015). Chapter Three - Positive Versus Negative Valence: Asymmetries in Attitude Formation and Generalization as Fundamental Individual Differences. In J. M. Olson & M. P. Zanna (Eds.), *Advances in Experimental Social Psychology* (Vol. 51, pp. 97–146). Academic Press. https://doi.org/10.1016/bs.aesp.2014.09.002

Gigerenzer, G. (2018). The Bias Bias in Behavioral Economics. *Review of Behavioral Economics*, *5*(3–4), 303–336.

Goldenberg, A., & Gross, J. J. (2020). Digital emotion contagion. *Trends in Cognitive Science*, online. https://doi.org/10.1017/CBO9781107415324.004

Goldenberg, A., Schöne, J., Huang, Z., Sweeny, T. D., Brady, T. F., Robinson, M. D., Levari, D. E., Zaki, J., & Gross, J. J. (2022). Amplification in the evaluation of multiple emotional expressions over time. *Nature Human Behavior*.

Goldenberg, A., Weisz, E., Sweeny, T., Cikara, M., & Gross, J. J. (2021). The crowd emotion amplification effect. *Psychological Science*, *32*(3), Article 3. https://doi.org/10.31219/osf.io/cn6qy

Goldenberg, A., & Willer, R. (2023). Amplification of emotion on social media. *Nature Human Behaviour*, *7*(6), Article 6. https://doi.org/10.1038/s41562-023-01604-x

Gorges, M., Porat, R., Schöne, J., & Goldenberg, A. (2024). *Aggregating Emotional Sequences Amplifies the Perception of Women as More Emotional Than Men*. https://doi.org/10.31219/osf.io/64dru

Gosling, S. D., Rentfrow, P. J., & Swann, W. B. (2003). A very brief measure of the Big-Five personality domains. *Journal of Research in Personality*, *37*(6), 504–528. https://doi.org/10.1016/S0092-6566(03)00046-1

Harris, R. B., & Paradice, D. (2007). *An Investigation of the Computer-mediated Communication of Emotions*. https://www.semanticscholar.org/paper/An-Investigation-of-the-Computer-mediated-of-Harris-Paradice/c891d020af8ec82660e4a50502626d9a3b29e6d4

Jackson, M. C. (2018). Eye gaze influences working memory for happy but not angry faces. *Cognition and Emotion*, *32*(4), Article 4. https://doi.org/10.1080/02699931.2017.1345720

Jackson, M. C., Wu, C. Y., Linden, D. E. J., & Raymond, J. E. (2009). Enhanced visual short-term memory for angry faces. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(2), Article 2. https://doi.org/10.1037/a0013895

Kelly, C., & Sharot, T. (2023). *Knowledge-Seeking Reflects and Shapes Well-Being*. OSF. https://doi.org/10.31234/osf.io/yd6j5

Kensinger, E. A. (2020). Negative emotion enhances memory accuracy. *Memory as Prediction*, *16*(4), Article 4. https://doi.org/10.7551/mitpress/13543.003.0005

Kensinger, E. A., & Schacter, D. L. (2008). Memory and emotion. In *Handbook of emotions, 3rd ed* (pp. 601–617). The Guilford Press.

Kort-Butler, L. A., & Hartshorn, K. J. S. (2011). Watching the Detectives: Crime Programming, Fear of Crime, and Attitudes about the Criminal Justice System. *The Sociological Quarterly*, *52*(1), 36–55. https://doi.org/10.1111/j.1533-8525.2010.01191.x

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2013). lmerTest: Tests for random and fixed effects for linear mixed effect models (lmer objects of lme4 package). *R Package Version*, *2*(6), Article 6.

Lerner, J. S., Li, Y., Valdesolo, P., & Kassam, K. S. (2015). *Emotion and Decision Making*. https://doi.org/10.1146/annurev-psych-010213-115043

Lindström, B., Bellander, M., Schultner, D. T., Chang, A., Tobler, P. N., & Amodio, D. M. (2021). A computational reward learning account of social media engagement. *Nature Communications*, *12*(1), 1311. https://doi.org/10.1038/s41467-020-19607-x

Luminet IV, O., Bouts, P., Delie, F., Manstead, A. S. R., & Rimé, B. (2000). Social sharing of emotion following exposure to a negatively valenced situation. *Cognition & Emotion*, *14*(5), Article 5.

Ma, G. W. S., Schöne, J. P., & Parkinson, B. (2024). Social sharing of emotion during the collective crisis of COVID-19. *British Journal of Psychology*, *115*(4), 843–879. https://doi.org/10.1111/bjop.12729

Manstead, A. S. R., & Fischer, A. H. (2001). Social appraisal: The social world as object of and influence on appraisal processes. In K. R. Scherer, A. Schorr, & T. Johnstone (Eds.), *Emotion: Theory, methods, research* (pp. 221–232). Oxford University Press.

Mattick, R. P., & Clarke, J. C. (1998). Development and validation of measures of social phobia scrutiny fear and social interaction anxiety. *Behaviour Research and Therapy*, *36*(4), Article 4. https://doi.org/10.1016/S0005-7967(97)10031-6

Mello, V. O. de, Cheung, F., & Inzlicht, M. (2022). *Twitter (X) use predicts substantial changes in well-being, polarization, sense of belonging, and outrage*. OSF. https://doi.org/10.31234/osf.io/4x5em

Näsi, M., Räsänen, P., Hawdon, J., Holkeri, E., & Oksanen, A. (2015). Exposure to online hate material and social trust among Finnish youth. *Information Technology & People*, *28*(3), 607–622. https://doi.org/10.1108/ITP-09-2014-0198

Öhman, A., Flykt, A., & Esteves, F. (2001). Emotion drives attention: Detecting the snake in the grass. *Journal of Experimental Psychology: General*, *130*(3), 466–478. https://doi.org/10.1037/0096-3445.130.3.466

Parkinson, B., & Simons, G. (2009). Affecting Others: Social Appraisal and Emotion Contagion in Everyday Decision Making. *Personality and Social Psychology Bulletin*, *35*(8), 1071–1084. https://doi.org/10.1177/0146167209336611

Piccardi, T., Saveski, M., Jia, C., Hancock, J. T., Tsai, J. L., & Bernstein, M. (2024). *Social Media Algorithms Can Shape Affective Polarization via Exposure to Antidemocratic Attitudes and Partisan Animosity* (arXiv:2411.14652). arXiv. https://doi.org/10.48550/arXiv.2411.14652

Redelmeier, D. A., & Kahneman, D. (1996). Patients' memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures. *Pain*, *66*(1), Article 1. https://doi.org/10.1016/0304-3959(96)02994-6

Rimé, B. (2009). Emotion elicits the socialsharing of emotion: Theory and empirical review. *Emotion Review*, *1*(1), Article 1. https://doi.org/10.1177/1754073908097189

Robertson, C. E., del Rosario, K., & Van Bavel, J. J. (2024). Inside the Funhouse Mirror Factory: How Social Media Distorts Perceptions of Norms. *Current Opinion in Psychology*, 101918. https://doi.org/10.1016/j.copsyc.2024.101918

Robertson, C. E., Pröllochs, N., Schwarzenegger, K., Pärnamets, P., Van Bavel, J. J., & Feuerriegel, S. (2023). Negativity drives online news consumption. *Nature Human Behaviour*, *7*(5), Article 5. https://doi.org/10.1038/s41562-023-01538-4

Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, *5*(4), Article 4. https://doi.org/10.1207/S15327957PSPR0504_2

Schirmer, A., & Adolphs, R. (2017). Emotion perception from face, voice, and touch: Comparisons and convergence. *Trends in Cognitive Science*, *21*(3), Article 3. https://doi.org/10.1016/j.tics.2017.01.001

Schöne, J. P., Garcia, D., Parkinson, B., & Goldenberg, A. (2023). Negative expressions are shared more on Twitter for public figures than for ordinary users. *PNAS Nexus*, *2*(7), pgad219. https://doi.org/10.1093/pnasnexus/pgad219

Schöne, J., Parkinson, B., & Goldenberg, A. (2021). Negativity spreads more than positivity on Twitter after both positive and negative political. *Affective Science*.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, *5*(2), 207–232. https://doi.org/10.1016/0010-0285(73)90033-9

Van Boven, L., White, K., & Huber, M. (2009). Immediacy bias in emotion perception: Current emotions seem more intense than previous emotions. *Journal of Experimental Psychology: General*, *138*(3), 368–382. https://doi.org/10.1037/a0016074

Wade, J. (2017, September 13). *Social Media Week London 2017 Recap*. Smart Insights. https://www.smartinsights.com/social-media-marketing/social-media-week-london-2017-recap/

Zhang, H., & Qu, C. (2020). Emotional, especially negative microblogs are more popular on the web: Evidence from an fMRI study. *Brain Imaging and Behavior*, *14*(5), 1328–1338. https://doi.org/10.1007/s11682-018-9998-6