

[PLEASE CLICK HERE TO DOWNLOAD THE MOST RECENT VERSION](#)

Emotion- versus Reasoning-based Drivers of Misinformation Sharing:

A field experiment using text message courses in Kenya*

Susan Athey¹, Matias Cersosimo¹, Kristine Koutout¹ and Zelin Li²

¹Graduate School of Business, Stanford University

²Sloan School of Management, Massachusetts Institute of Technology

November 2022

*We would like to acknowledge the team at First Draft News, now Information Futures Lab, Claire Wardle, Shaydanay Urbani, and Rory Smith for developing the text message courses in this study. We would also like to acknowledge Tommy Shane and Laura Garcia for their contributions to early discussions of the project. We gratefully acknowledge the Golub Capital Social Impact Lab and the Africa Infodemic Response Alliance at the World Health Organization for their financial support of this study. This experiment was preregistered with the AEA RCT Registry under AEARCTR-0009721 and reviewed by the Stanford Institutional Review Board under protocol 57430.

Abstract

Two leading hypotheses for why individuals unintentionally share misinformation are that 1) they are unable to recognize that a post contains misinformation, and 2) they make impulsive, emotional sharing decisions without thinking about whether a post contains misinformation. Much of the literature on interventions to counter misinformation focuses on the first hypothesis and tests interventions that educate social media users about reasoning-based techniques employed in social media posts to mislead them; however, other work shows that emotions are important for the spread of and belief in misinformation, supporting a focus on the second hypothesis and emotion-based techniques. To learn whether interventions to counter reasoning- or emotion-based techniques are more effective, or whether the approaches are complementary, we evaluate three distinct versions of a five-day, low-cost, and scalable text message educational course in a field experiment with approximately 9,000 participants in Kenya. We assess the impact of the courses through a pre-post survey design eliciting intentions to share, finding that all treatment courses work, decreasing misinformation sharing 28% on average relative to no text message course. Treating the emotional drivers of misinformation sharing in the “Emotions” course is more effective than teaching about reasoning-based techniques either alone in the “Reasoning” course or in combination with emotion-based techniques in the “Combo” course. Moreover, the Emotions course performs best on misinformation posts that use emotion-based techniques, and does no worse than the Reasoning or Combo courses on misinformation posts that use reasoning-based techniques. In a follow-up experiment approximately two months later, 88% of the treatment effect of the three courses on misinformation sharing persists.

1 Introduction

The growth and pervasiveness of misinformation in recent years has shifted attention from debunking strategies that intervene one piece of misinformation at a time, to technique-based interventions that educate social media users about misinformation broadly. These interventions mostly focus on reasoning-based techniques to mislead users in social media posts like false dichotomies or out-of-context images, but some also include emotion-based techniques like the use of emotion- and moral-laden language in posts.¹ These emotion-based techniques, while less studied in the intervention literature, have been shown to be important for the spread of misinformation (Brady et al., 2017; Pröllochs et al., 2021) and in how likely people are to believe in misinformation (Martel et al., 2020). Another important differentiator between these two types of techniques is that strategies to counter each interfere at different stages in the user’s decision-making process.

Consider a user who sees a post in their news feed they want to share. In the first stage of their decision-making process, the user decides whether to share the post immediately or to evaluate if the post contains misinformation. Emotional language in the post can trigger an impulsive sharing decision at this stage, so strategies to counter emotion-based techniques aim to tame this impulse. Conditional on not sharing in the first stage, the user evaluates whether the post contains misinformation in the second stage and subsequently make an informed sharing decision. Strategies to counter reasoning-based techniques aim to provide the user with the tools and strategies to evaluate posts in this second stage. While either strategy might be effective alone, because each

¹See Compton et al. (2021); Lewandowsky and Van Der Linden (2021); van der Linden (2022) for reviews.

intervenes at a different stage in the user’s decision-making process, the two strategies may be particularly powerful in combination. By first determining whether it is more effective to teach reasoning- or emotion-based techniques, and then measuring the impact of teaching both, we can learn about the motivations driving users to share misinformation and apply this knowledge to refine future interventions.

To study these questions, we collaborated with experts at the nonprofit organization First Draft News to modify a text message course aimed at countering misinformation sharing that they developed for distribution through governments and nongovernmental organizations. We tested this low-cost, five-day text message course distributed via MMS or WhatsApp in a field experiment by recruiting Facebook users in Kenya, a likely target for the scaled intervention. Text messages as a delivery mechanism for educational content break access barriers, particularly in a lower-income context, for people who cannot come to synchronous classes, do not have access to a personal computer with steady internet service to watch video lectures online, or have limited data on their mobile phones for otherwise engaging with videos or education platforms. By adding a pre-post online survey to the text message course, we evaluate the course using intentions to share. We follow up with participants in a survey administered seven to eleven weeks after completing the course to measure long-run effects.

We test three distinct versions of the text message course that teach users to counter either (i) *reasoning*-based techniques, (ii) *emotion*-based techniques, or (iii) a combination of the two. For example, the “Emotions” course teaches participants how misinformation creators use emotional language to manipulate users, and prompts them to stop and question a post that elicits a strong emotional reaction. Participants are randomly assigned to receive one of the three treatment courses—Emotions, “Reasoning,” or “Combo”—or one of two baselines. The “No-course” baseline receives no course between the pre- and post-survey. The “Facts” baseline is exposed to daily facts about misinformation but no technique-based education. In addition to randomizing participants into one of five text message course intervention assignment groups, we independently randomize participants to see a “high-dosage” version of an accuracy nudge (Pennycook et al., 2021) in the pre- and post-survey, where participants rate the accuracy of a post before making a sharing decision.

The survey elicits sharing decisions for non-misinformation posts and three types of misinformation posts corresponding to some techniques taught in each course: Reasoning posts contain reasoning-based techniques; Emotions posts contain emotion-based techniques; and Combo posts include both. By examining heterogeneity in misinformation sharing by type of post, we can observe the cross-technique impact of courses on sharing different types of misinformation.

We find that all treatment courses effectively and persistently reduce misinformation sharing, but the Emotions course is most effective. Perhaps surprisingly, there is no incremental value to teaching about reasoning-based techniques when participants are also exposed to strategies to counter emotion-based techniques. Importantly, the Emotions course not only does better on posts with emotion-based techniques but does not do meaningfully worse than either the Reasoning or Combo courses on posts with reasoning-based techniques only.

On average, the treatment courses reduce misinformation sharing by 18 percentage points (p.p.) immediately upon completion of the course, which is 28% of the 64% sharing rate in the group not exposed to a text message course. The treatment courses also reduce misinformation sharing

by 10 p.p. relative to the Facts baseline, showing that the effect is not driven by the course making misinformation salient each day. The effect of the treatment courses persists, reducing misinformation sharing 9 p.p. more than the Facts baseline after approximately two months, or 88% of the short-run effect. In addition, the accuracy nudge works as expected, reducing misinformation sharing by 7 p.p., but not as well as the treatment courses (18 p.p.). We find no evidence that combining the accuracy nudge with the treatment courses has any incremental effects, suggesting that any complementarity is small at best.

Our finding that the Emotions course is consistently the most effective suggests that people make impulsive, emotions-driven sharing decisions that lead to misinformation sharing. This finding is concerning for a number of reasons. First, heightened emotional states are correlated with greater belief in misinformation [Martel et al. \(2020\)](#). Second, the misinformation literature to date focuses primarily on evaluating interventions that counter reasoning-based techniques, like debunking interventions and technique-based interventions to determine whether a post is misinformation, rather than identifying or developing interventions that counter emotion-based techniques. Last, posts that contain emotional language are dispersed more broadly and quickly than posts without emotional language based on text analyses of social media content ([Brady et al., 2017](#); [Pröllochs et al., 2021](#)). The importance of this last point is further underscored by our results showing that posts with emotional language are precisely the types of posts for which the Emotions course has the greatest advantage over the other courses. So, not only is the Emotions course the most effective overall, it has the greatest advantage in changing behavior on the types of posts that are most concerning for promoting the belief in and the spread of misinformation.

One way to understand the psychological mechanisms driving the success of the Emotions course is through the influential Elaboration Likelihood Model (ELM) of persuasion ([Petty and Cacioppo, 1986](#)) in the psychology literature. The ELM posits that people use either central route processing, a high level of thinking (elaboration) focused on the merits of a persuasive argument, or peripheral route processing, a low level of thinking that uses environmental cues and heuristics, to evaluate a persuasive argument (or some combination) ([Petty and Hineskamp, 2017](#)). One finding in this body of literature based on lab studies is that emotions can affect how much central versus peripheral elaboration a person does when evaluating a persuasive argument and, when peripheral route processing dominates, whether the person ultimately accepts or rejects the argument (see [Petty and Briñol \(2015\)](#) for a review).

The primary technique taught in the Emotions course is to “Stop and Question” content that elicits a strong emotional reaction, which is a type of emotion regulation through “cognitive reappraisal.” The emotions regulation literature in psychology finds in lab experiments and observational studies that cognitive reappraisal is frequently successful at changing perceived emotions in the short-run through changing how a person thinks about the emotions they are experiencing. Cognitive reappraisal also is more successful in the long-run than other emotion regulations strategies like suppression (see [McRae and Gross \(2020\)](#) for a review). Thus, one way to understand how the Emotions course might work psychologically is it teaches participants that misinformation often uses emotions to manipulate them (e.g., by triggering peripheral instead of central route processing). Then, the Emotions course provides a strategy participants can use to “reappraise” how they think about those emotions when they experience them. This reappraisal could trigger a central route process that leads to higher elaboration and could also lead participants to evaluate a post based on the merits of its argument rather than heuristics.

Our experiment also evaluates another type of misinformation intervention that has been broadly influential: an “accuracy nudge” (Pennycook et al., 2021). The accuracy nudge serves two purposes in our study. First, the accuracy nudge captures features of each of the treatment courses, but it does not include educational content. Like the “Stop and Question” strategy in the Emotions course, the accuracy nudge we implement creates distance in time from exposure to the post, with any accompanying emotional reaction, to the sharing decision. But like the strategies in the Reasoning course, the accuracy nudge encourages users to think about the accuracy of the post. So, by comparing the performance of the treatment courses to the accuracy nudge, we can evaluate whether the more in-depth educational content adds value beyond these features of the courses. Second, accuracy nudges have been evaluated using on-platform behavior as an outcome (Pennycook et al., 2021), while our study relies on survey outcomes. By including the accuracy nudge in our experiment and measuring its effect on survey outcomes, we can make internally valid comparisons to this externally validated intervention.

The inclusion of an accuracy nudge in the pre- and post-surveys is subtle, namely because the intervention itself is a survey question. Pennycook et al. (2021) implement an accuracy nudge by sending an unsolicited direct message to a subset of Twitter users who had recently engaged with untrustworthy news sites that contained one post with a question: “How accurate is this headline?” Users did not need to engage with the question; instead, Pennycook et al. estimate an intent to treat on exposed users and find increased discernment of misinformation in the 24 hours following exposure. Offer-Westort et al. (2021) include a similar accuracy nudge in their experiment—one post with a question about the post’s accuracy. Our experiment evaluates the impact of a high-dosage version of an accuracy nudge, in which participants who are randomly assigned to see the accuracy nudge score the accuracy of *each* post before they make the sharing decision for that post. Those participants *not* assigned to see the accuracy nudge are instead first asked about their intention to share on a set of posts.

When interpreting the findings from our experiment, it is important to consider whether our results are externally valid; that is, whether they reflect the treatment effects that could be expected if the course was scaled up. First, we can contemplate the population to whom a “scaled up” version of this intervention might be applied.² The method of attracting users to take the course would affect the population that received it. One approach might be to conduct a public information campaign using advertisements or public service announcements on social media, tailored to a country or region where misinformation was viewed as a challenge. As our experiment recruits participants using advertisements on social media in Kenya where the prevalence of misinformation has been particularly concerning,³ we draw a similar population to the one that might be targeted in a scaled-up intervention. However, our campaign attracted only about 9,000 users to complete the study and paid users with mobile airtime, both of which contribute to selection. If the campaign was expanded, a different selection of users might be attracted, and attrition patterns might also be different.

In terms of what has been called “naturalness,” in many dimensions our experiment mimics closely to what might be contemplated in a scaled-up intervention. In both our study and a potential large-

²We reference the SANS (Selection, Attrition, Naturalness, Scaling) transparency conditions in List (2020) in our discussion of external validity.

³The Africa Infodemic Response Alliance at the World Health Organization has highlighted the prevalence of misinformation on social media in Kenya (Nguyen and Cecchini, 2021)

scale implementation, participants would enter their phone number (on a web page as opposed to in a survey) to receive daily text messages via a messaging app (our study gave the choice of MMS or WhatsApp) for five days. Offering explicit incentives, as in our experiment, would increase take-up of the course if implemented at a larger scale, particularly among those who may not think they need an intervention on misinformation sharing. A school, government body, or company might also require or incentivize (e.g., with prizes or lotteries) individuals to participate in the text message course.

Arguably the most important limitation to the external validity of our approach is that we measure outcomes through surveys. Although this is the most common measurement method in the literature to date (with the notable exception of [Pennycook et al. \(2021\)](#)), survey measures also have known limitations. For example, survey measures for misinformation sharing outcomes may be subject to experimenter demand effects and other biases due to the hypothetical nature of the questions. We address this shortcoming in three ways. First, as discussed above, we use the accuracy nudge as an internally valid benchmark to an externally validated intervention. Second, our survey measures are separated in time and context from the text message courses. Further, we measure long-run outcomes to evaluate the consistency and duration of results, as experimenter demand effects are likely to be most pronounced directly after interacting with the experiment and the associated course. Third, we use qualitative questions in the follow-up survey to determine whether participants *report* implementing techniques from the course in their social media behavior and whether treated participants do so more than the baseline. The performance of the courses relative to the accuracy nudge, the persistence of the effect of the treatment courses, and the reported behavior in the follow-up survey all support the conclusion that educating users about emotion-based techniques is more effective than the alternatives or baselines we test.

Our study is the first to our knowledge to evaluate treating reasoning- versus emotion-based drivers of misinformation sharing both separately and in combination. The experiment has sufficient power to compare the two versions of the course targeting different mechanisms, as well as their interaction, such that we can reject the hypothesis that the course teaching both emotion and reasoning-based techniques is more effective than a course focused on emotion-based techniques alone. We test a novel intervention, the first text message course for treating misinformation sharing to our knowledge.^{4, 5} Even though online misinformation is a global phenomenon, we also conduct one of the first evaluations of misinformation interventions outside of the U.S. and high-income countries (with the notable exception of [Offer-Westort et al. \(2021\)](#)). Other studies that use survey experiments recruit participants globally ([Arechar et al., 2022](#); [Basol et al., 2021](#); [Roozenbeek and Van Der Linden, 2019](#); [Roozenbeek et al., 2022](#)) but do not focus on low-income countries in particular. We next provide context for these contributions in the related literature.

⁴[Offer-Westort et al. \(2021\)](#) send text messages through a chatbot and all messages are delivered in the same session.

⁵Text message courses have proven to be useful and effective in stimulating positive behavioral changes in previous studies ([Armanasco et al., 2017](#)).⁶

1.1 Related Literature

The text message course we test in this study is categorized as an “inoculation” intervention. Inoculation theory can be traced back to as early as the 1960s when [McGuire \(1961\)](#) developed the theory to understand persuasion-resistance with an analogy similar to that of a vaccine. The theory has only recently been utilized in the context of misinformation, where the first set of literature primarily studied inoculation in the context of “prebunking” climate change misinformation ([Cook et al., 2017](#); [Van der Linden et al., 2017](#)). This study contributes to the second wave of inoculation interventions focused on technique-based approaches, as opposed to “issue-based” approaches that treat one piece of misinformation at a time ([Lewandowsky and Van Der Linden, 2021](#)).

The most relevant work to this paper evaluates one of two interventions that teach both reasoning- and emotion-based techniques used to mislead users in online posts: the Bad News game⁷ and the Inoculation Science videos.⁸ Bad News is an online game in which players role-play as a fake news creator attempting to amass followers while maintaining credibility. The game takes approximately 15 minutes to complete by earning six badges corresponding to different techniques used to spread misinformation. One of the badges is for emotional language, while others focus on reasoning-based techniques like impersonation and conspiracy theories. The video series created by Inoculation Science similarly includes a video on emotional manipulation and five other short videos covering reasoning-based techniques like false dichotomies or ad hominem attacks.

In a series of papers starting with [Roozenbeek and Van Der Linden \(2019\)](#), the Bad News game has been shown to decrease susceptibility to misinformation ([Basol et al., 2020](#); [Roozenbeek et al., 2021](#)), including up to 13 weeks after playing the game with regular reminders ([Maertens et al., 2021](#)). The general setup in these papers is that a pre- and post-survey embedded in the game asks participants to rate the reliability, manipulativeness, and/or accuracy of posts. Then, they use a convenience sample of people who play the game online and complete the surveys to evaluate change in survey outcomes before and after playing the game. [Basol et al. \(2020\)](#) uses Prolific instead of a convenience sample so they can evaluate the performance of the game relative to a control group that plays Tetris.

In [Roozenbeek et al. \(2022\)](#), the authors evaluate each Inoculation Science video on technique recognition, confidence, trustworthiness, and sharing discernment. Participants on Prolific view one of the six videos, then [Roozenbeek and Van Der Linden](#) show them posts that either do or do not include the technique from the relevant video (i.e., the emotional manipulation video is evaluated using posts that either include emotional manipulation or are non-misinformation). [Roozenbeek et al. \(2022\)](#) find that all videos improve most outcomes, but they do not explicitly discuss the differences they observe between the videos, likely because there are differences between the posts used to evaluate each video. Their results do not show that the emotional manipulation video does best in any of the outcomes they assess (including willingness to share, the metric we use in this paper, where they find no differences between the videos). [Roozenbeek et al.](#) also evaluate two of the videos on YouTube (emotional manipulation and false dichotomies) by randomly showing one of two of the videos to users as ads on YouTube. A subset of those users completed a one-question survey about whether the technique they learned about in the video was present in the headline

⁷<https://www.getbadnews.com/books/english/>

⁸<https://inoculation.science/inoculation-videos/>

they were shown. Compared to a control who did not see the videos, users were more likely to correctly identify the technique.

In another study with similarities to our own, [Offer-Westort et al. \(2021\)](#) conduct an adaptive experiment in sub-Saharan Africa (Kenya and Nigeria) to evaluate 11 interventions and determine both the best treatment overall and heterogeneity in optimal treatment assignment. Like our study, [Offer-Westort et al.](#) recruits participants through Facebook ads, but the interventions they test are deployed through a chatbot. [Offer-Westort et al.](#) also use a pre-post survey design, but because the chatbot delivers the intervention instantaneously, they measure intentions to share in the same session in which the intervention is delivered. The “Facebook tips” intervention that [Offer-Westort et al.](#) test, in which the chatbot provides users with strategies to counter misinformation, is similar to a text message course.⁹ Our intervention differs from theirs in that our course is delivered over multiple days (compared to in the same session) through either text or WhatsApp (as opposed to Facebook Messenger) and contains some interactive questions.

Outside of the misinformation context, educating participants about how an organization might be manipulating them has been employed in other contexts. For example, [Bryan et al. \(2016\)](#) use a randomized experiment to evaluate an intervention that teaches adolescents about how the food industry manipulates consumers; it was shown to improve the diets of the participants relative to several baselines. [Bryan et al. \(2016\)](#) interprets the finding as evidence that adolescent motivation around autonomy provides a symbolic reward (“feeling like a... respect-worthy person acting in accordance with important values shared with peers”) that competes with the short-term temptation of unhealthy eating, a theory that has grounding in the neuroscience literature ([Telzer et al., 2014](#)). In our study, the symbolic reward around avoiding manipulation by an organization might compete against the temptation of succumbing to fear, shock, or outrage and sharing a post.

2 Experimental Design

The goal of this experiment is to study the effect of different text message courses on misinformation sharing behavior. Figure 1 diagrams the experimental design. We recruited participants through Facebook ads to complete a five-day text message course, plus a pre- and post-survey, for a mobile airtime payment. Participants who clicked an ad were directed to the pre-survey on Qualtrics, where they were first randomized on whether they would see an accuracy nudge in the pre- and post-survey. Conditional on completing the pre-survey, participants were randomized into one of the text message course interventions and enrolled in the course. Participants in all course interventions, except the No-course baseline, received one text message a day for five days starting on the day they completed the pre-survey.¹⁰ On the last day of the course, participants received a link to complete the post-survey on Qualtrics. Participants who completed the post-survey were paid KSH 500 (about \$4 in U.S. dollars) in mobile airtime. Seven to eleven weeks later, participants who completed the post-survey were randomized on whether they would see a prime in the text message recruiting them to a follow-up survey. Participants who completed the follow-up survey were paid KSH 350 (about \$3 in U.S. dollars).

⁹Facebook tips were also developed by First Draft News.

¹⁰Participants in the No-course baseline received the Combo course after the post-survey.

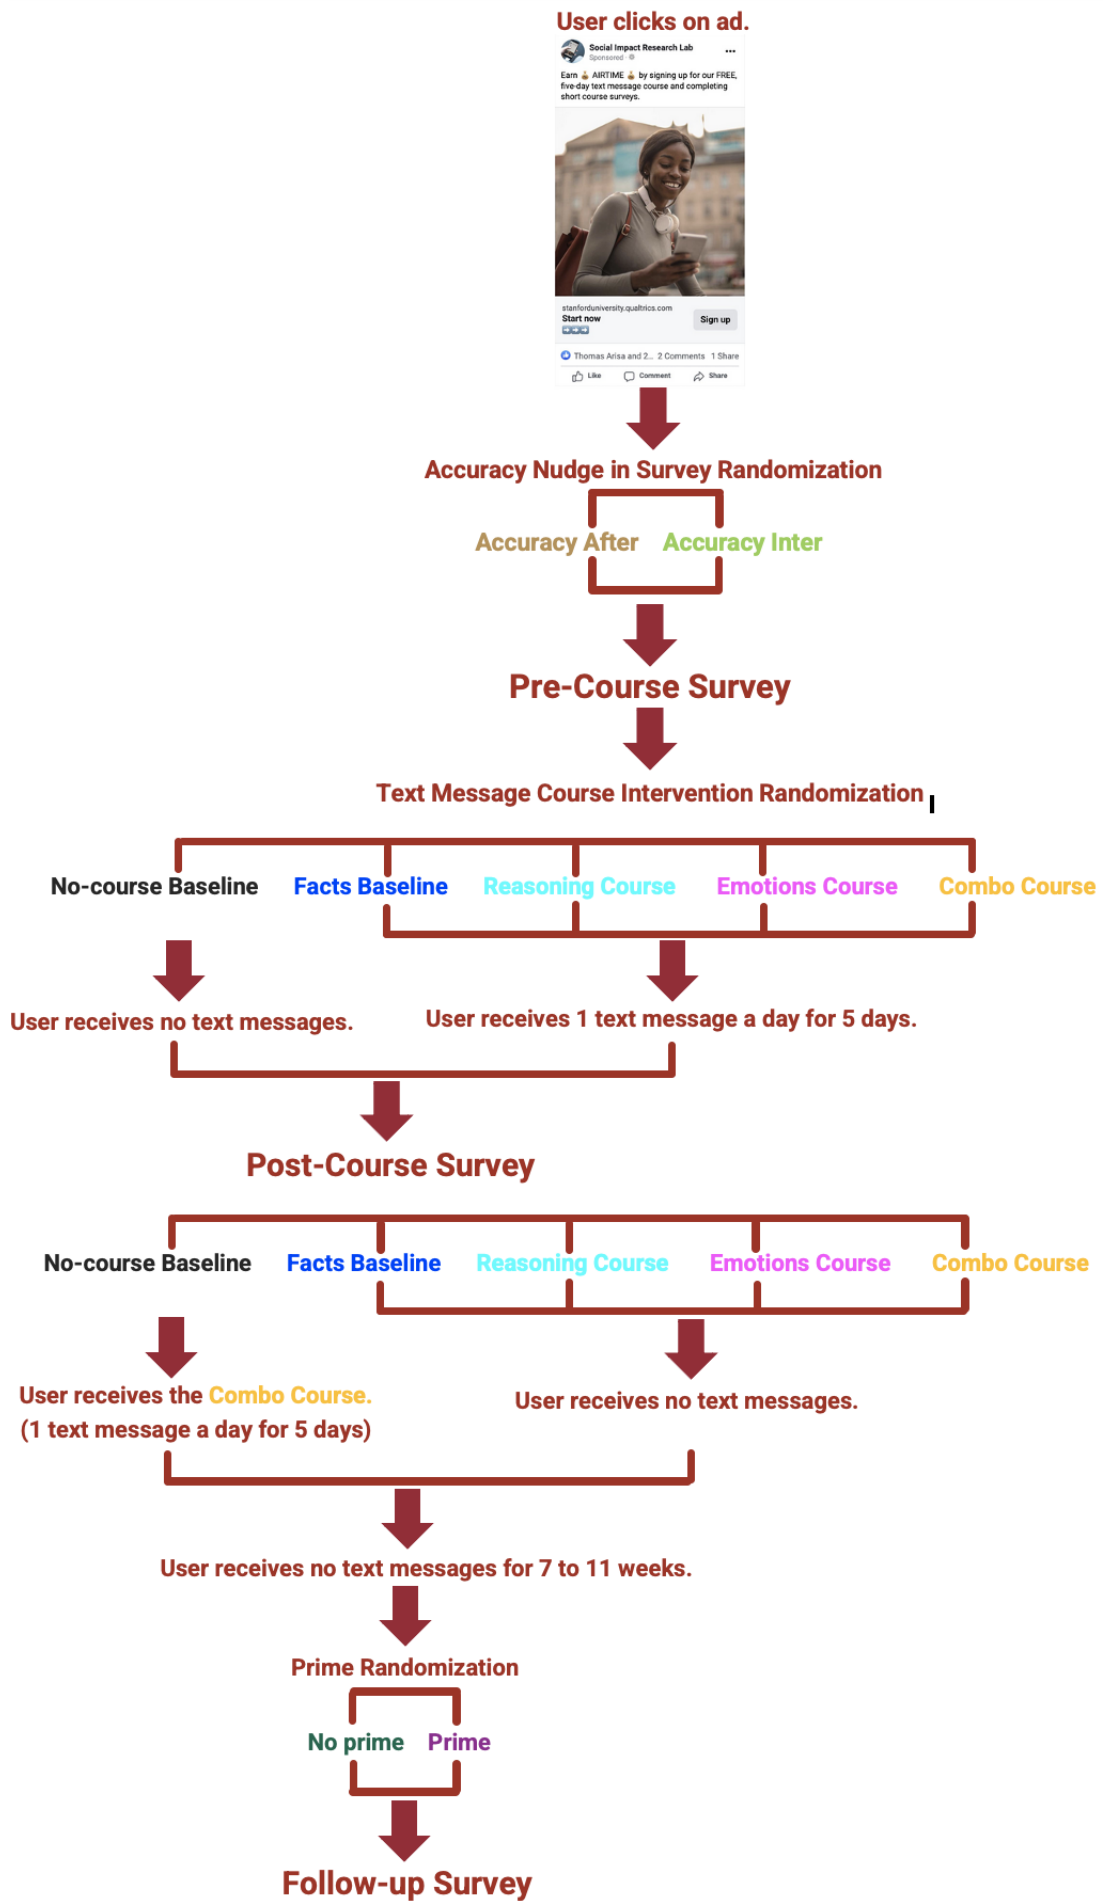


Figure 1: Experiment Design

In this section, we first detail the Stage 1: Main Experiment that encompasses recruitment to the pre-survey through the post-survey. There are two independent randomizations in this experiment: the accuracy nudge in the surveys and the text message course interventions. There are five treatment course interventions; hence, the experiment is a five (course interventions) by two (accuracy nudge) factorial design. For the majority of our analysis, we aggregate over exposure to the accuracy nudge. We exploit the factorial design in Results section 3.5 to learn whether the treatment courses are more or less effective than the accuracy nudge at decreasing misinformation sharing, and whether there are incremental effects of combining the accuracy nudge with the treatment courses.

Next, we detail the Stage 2: Follow-up Experiment, in which there is a third independent randomization into whether participants are exposed to a prime. The follow-up survey contains no accuracy questions, so there is no accuracy nudge for any group. Lastly, we discuss implementation of the experiments.

2.1 Stage 1: Main Experiment

We recruited our sample from English-speaking Facebook users in Kenya who were at least 18 years old.¹¹ Ads ran on Facebook and Messenger, where users saw the opportunity to “Earn AIRTIME by signing up for our FREE, five-day text message course and completing short course surveys.” Clicking on an ad led participants to the pre-survey. Participants were told that those who completed the pre-survey, text message course, and the post-survey would receive KSH 500 in airtime, conditional on correctly answering all attention checks (example of an attention check question is included in Online Appendix D).¹² Appendix Table B1 shows our participant funnel from impressions to completion, with associated ad costs (\$1.21 for post-survey completion only and about \$2 for follow-up survey completion too). Online Appendix A contains a sample of the ads we ran.

Participants enrolled in the course at the end of the pre-survey by providing their phone number and choosing whether to receive the text message course via MMS or WhatsApp, a messaging platform with 97% penetration in Kenya.¹³ Upon enrollment, participants received the first of five daily text messages delivered at the same time (the time of enrolment) each day, unless they were assigned to the No-course baseline. Each day’s message was equal to or less than 1200 characters long. We next describe the five text message course interventions.

2.1.1 Text Message Course Interventions

First Draft News, a nonprofit organization focused on protecting communities from misinformation, has been a leader in developing interventions to combat misinformation sharing. For example, First Draft developed what is likely the most widely distributed intervention to date: Facebook’s

¹¹Kenya was selected by the Africa Infodemic Response Alliance in the World Health Organization due to the prevalence of misinformation on Kenyan social media.

¹²We paid KSH 55 to participants who did not correctly answer all attention checks to compensate them for their time, but participants were not told in advance we would do so.

¹³87% of participants chose to receive the course via WhatsApp.

“Tips to Spot False News.”¹⁴ Following the announcement of an “infodemic” by the World Health Organization (WHO) in 2020, First Draft developed text message courses as a low-cost intervention for governments and nongovernmental organizations to distribute. In addition to being relatively low-cost to develop and distribute, the text message course delivery mechanism makes the content more accessible to users in lower-income countries than online games and videos, which can be costly to engage with because they use data on mobile phones. We collaborated with First Draft to develop three distinct versions of this text message course to test two approaches to decreasing misinformation sharing and their interaction.

First, the Emotions course is motivated by the idea that people may not stop to think about whether a post is misinformation before sharing it because they make an emotions-driven decision. The course focuses on teaching a “Stop and Question” strategy to counteract emotion-based techniques used in social media posts to exploit people’s intuitive reaction to emotional stimulation. In particular, the course teaches that evoking emotions like fear, anger, and superiority are techniques often used to induce users to make an emotions-driven sharing decision.

Second, the Reasoning course is motivated by the idea that people may share misinformation because they *do not know* a post contains misinformation. The course focuses on teaching reasoning-based techniques used in social media posts to exploit a user’s lack of knowledge about context, sources, and other information, then teaching strategies to evaluate whether a post contains misinformation. Specific techniques taught by the course concern misleading graphs, imposter websites, and eyewitness media.

Each of these courses target a different point in the decision-making process in which a user first decides whether to evaluate a post or share it immediately; conditional on not sharing in this first stage, the user employs the strategies at their disposal to evaluate the post and then make a sharing decision. The Emotions course focuses on inducing people to pause in the first stage to evaluate whether a post is misinformation, which is a necessary precondition for the user to apply the strategies learned in the Reasoning course to evaluate the post. On the other hand, pausing to evaluate whether a post is misinformation is irrelevant if the user is unable to evaluate the post and distinguish non-misinformation from misinformation. In other words, the strategies taught in both courses may be necessary to have the greatest impact on misinformation sharing.

The third Combo course teaches the concepts from both the Emotions course and the Reasoning course to test if these approaches are complementary. To ensure that both the length of the text messages and the number of days in the course are the same as the Emotions and Reasoning courses, the content of those courses was streamlined.

In addition to the three treatment courses, there were two baselines. In the No-course baseline, participants received no text messages between the pre-survey and the post-survey. Conditional on completing the post-course survey, participants in the No-course baseline received the Combo course.¹⁵ The Facts baseline sent a text message each day with a fact about misinformation, such as “During the Roman Empire...leaders used misinformation to come to power.” The purpose of

¹⁴Guess et al. (2020) find that these tips, which were promoted at the top of Facebook users’ news feed in 14 countries in April 2017, improved misinformation discernment in a survey experiment similar to ours.

¹⁵As participants in the No-course baseline were paid upon completing the post-survey, they received the course but were not paid to engage with it.

the Facts baseline was to separately identify the effect of the educational content in the course and the potential salience effect from a daily text message on misinformation. All courses are provided in Online Appendix B.

2.1.2 Survey Structure

This section details the pre-survey, which participants accessed by clicking on a Facebook ad that linked them to the survey on Qualtrics, and the post-survey, which participants accessed by clicking on a link in the last text message they received as a part of the text message course intervention. Figure 2 shows the order of questions for each of the pre- and post-surveys. In the pre-survey, participants responded to social media posts, gave reasons for sharing specific posts, provided demographic information, and enrolled in a “free text message course that teaches you how to protect against misinformation.” In the post-survey, participants only responded to social media posts and again gave reasons for sharing specific posts. Online Appendix C contains screenshots of the pre-survey.¹⁶



Figure 2: Pre- and Post-survey Structure

Both the pre-survey and the post-survey showed participants ten social media posts and asked them two questions for each post:

1. Would you share this post? (Binary yes/no)
2. To the best of your knowledge, how accurate is the claim in the above post? (4-point Likert scale)

¹⁶The post-survey asked the sharing and accuracy questions in the same manner as the pre-survey, so screenshots of the post-survey provide no new information.

The accuracy nudge randomization that determined the order in which those questions were asked is detailed towards the end of Subsection 2.1.2.

At the end of the 20-question series (ten posts times two questions), we asked two open-ended free text questions to collect information about participants' reported reasons for their sharing decisions. Participants were shown a post that they previously reported wanting to (not) share, then asked "What made you (NOT) want to share it?" If they did not intend to share any of the posts, they did not see the free-text question on sharing. Similarly, if they always reported wanting to share, they did not see the free-text question on not sharing.

Participants saw one social media post used as an attention check and nine other posts, which were drawn from 62 social media posts we created for the pre- and post-surveys.¹⁷ To create 60 of the posts, we started from 15 health-related facts, verified by a fact-checking organization or published in an academic paper. For each fact, we created four posts. In addition, we created two attention check posts, randomizing which post participants saw in the pre- versus the post-survey. The four posts we created based on each fact were a:

- Baseline non-misinformation post ("non-misinformation")
- Misinformation post using an emotion-based technique ("Emotions")
- Misinformation post using a reasoning-based technique ("Reasoning")
- Misinformation post using both emotion- and reasoning-based techniques ("Combo")

We developed these four versions so that we could observe a participant's behavior when the same information was presented with and without manipulation. The correspondence of the types of misinformation to the treatment courses allows us to learn whether the treatment courses change behavior only on posts that employ the techniques taught therein or whether they have cross-technique effectiveness. If the Emotions course only changes behavior on Emotions posts, and similarly the Reasoning course only changes behavior on Reasoning posts, the policy implications would differ than if one course is as good or better on all types of misinformation (as we find).

In both the pre- and post-survey, each participant saw three non-misinformation posts, two Emotions posts, two Reasoning posts, two Combo posts, and one attention check post, in random order, for a total of ten posts. Importantly, the three non-misinformation posts a participant saw in the pre-survey were matched based on the "fact" to one Emotions, one Reasoning, and one Combo misinformation post the participant saw in the post-survey. We use these three non-misinformation posts in the pre-survey and their corresponding misinformation posts in the post-survey to construct our outcome. Each of the *other* posts shown to participants in the pre- and post-survey were randomly drawn from one of the remaining 12 facts, without replacement. Subsection 2.1.2 details the development of these posts.

The idea for this outcome measure is that the non-misinformation posts allow us to observe a participant's baseline propensity to share a particular fact in the pre-survey. Then, we can see how each

¹⁷We created an additional two non-misinformation posts and two attention check posts for the follow-up survey.

Fact	Pre	Post	Corr/Non-corr
1	non-misinformation	Emotions	Corr
2	non-misinformation	Reasoning	Corr
3	non-misinformation	Combo	Corr
4	None	Emotions	Non-corr
5	None	Reasoning	Non-corr
6	None	Combo	Non-corr
7	None	non-misinformation	Non-corr
8	None	non-misinformation	Non-corr
9	None	non-misinformation	Non-corr
10	Emotions	None	Non-corr
11	Emotions	None	Non-corr
12	Reasoning	None	Non-corr
13	Reasoning	None	Non-corr
14	Combo	None	Non-corr
15	Combo	None	Non-corr

Table 1: Example of Posts for each Fact, by Pre-Survey and Post-Survey

Notes: The first column indicates the fact number. The second and third columns indicate an example random draw of posts to show a participants in the pre-survey and post-survey, respectively. The fourth column shows for which facts the example participant sees a non-misinformation post in the pre-survey and a corresponding misinformation post in the post-survey about the same fact.

intervention affects the participant’s propensity to share the same fact presented as misinformation (i.e., we are controlling for a participant’s interest in sharing posts about a specific topic). So, when we observe that a participant reports does not intend to share a misinformation post about autism in the post-survey, for example, our primary outcome measure rules out the possibility that the participant is not interested in posts about autism generally. We also test alternative outcome measures in Results section 3.3 and find that our results are robust.

Social Media Posts The posts participants see in the surveys were designed to have the same information in all four versions of a post based on one fact. We kept the social media poster and their profile picture the same for all versions of a post, although these elements changed from fact to fact. We varied other characteristics between posts and across facts. Specifically, we varied the length and tone of the posts, and the presence or absence of an image, link, and/or hashtags on the posts. The heterogeneity in the posts was designed to make the detection of misinformation nontrivial. For example, we did not want participants to note that exclamation points always appeared in misinformation but not in non-misinformation posts. While exclamation points may be correlated with misinformation, they do not make a post misinformation.

The three types of misinformation posts used techniques taught in the three treatment courses described in section 2.1.1, so that we could learn whether the treatment courses changed behavior only

on posts that employed the techniques taught therein or had cross-technique effectiveness. Emotions posts used emotional and moral-laden language. We targeted emotions taught in the course (fear, anger, and superiority), as well as other emotions such as making people feel like something *must* be wrong. Reasoning posts used techniques from the course and from the FLICC taxonomy (Fake expert, Logical fallacies, Impossible expectations, Cherry picking, Conspiracy theories) developed by Cook (2020). We targeted concepts not taught in the course, as well as those taught in the course, to avoid “teaching to the test.” Online Appendix D contains examples of the posts used in the survey.

We aimed to make the posts realistic representations of what social media users would see in Kenya. Posters were designed to be media reporters and other individuals who would not be in a user’s network. Social media users would see these types of posts through groups on Facebook or public forums such as Twitter. To verify that the posts were representative of what Kenyans see on the internet, we recruited Kenyan workers on Upwork to evaluate the content of our posts, and we made changes based on their suggestions.

Accuracy Nudge Participants were randomly assigned to see the accuracy nudge at the beginning of the pre-survey. This treatment assignment determined the *order* of the first 20 questions, which elicited sharing intentions and accuracy for each of ten posts in both the pre- and post-survey. Figure 2 shows the question order for those who were, and were not, assigned to see the accuracy nudge. Participants treated with the accuracy nudge (Accuracy Inter group) were first asked about the accuracy of a post and then were asked to make a sharing decision about the same post, for each of the ten posts (i.e., questions interweave). Participants not treated with the accuracy nudge (Accuracy After group) were asked to make a sharing decision about all ten posts and then were asked about the accuracy of all ten posts.

Accuracy nudges have been shown to be an effective means of decreasing misinformation sharing in laboratory experiments, survey experiments, and on Twitter (Offer-Westort et al., 2021; Pennycook et al., 2021). The mechanism through which accuracy nudges aim to treat misinformation sharing – priming a person to think about the accuracy of the post – has synergies with strategies in both the Emotions and Reasoning courses. The accuracy nudge forces users to pause before making the sharing decision, a strategy the Emotions course encourages users to employ. The accuracy nudge also encourages users to evaluate the accuracy of posts before making a sharing decision, for which the Reasoning course provides strategies.

This survey design element allows us to achieve two goals. First, we can learn whether the text message course provides educational value beyond encouraging users to pause and evaluate posts before sharing them. For example, in the Emotions course, the “Stop and Question” strategy is motivated by describing how emotional language is used to mislead users and manipulate them into sharing. This educational content could increase participants’ desire to avoid sharing misinformation, and so the accuracy nudge allows us to evaluate whether that content has added value. Second, as an externally-validated intervention, the accuracy nudge provides an internally-valid benchmark for the text message courses.

2.2 Stage 2: Follow-up Experiment

The experimental design to this point was prespecified. When we learned that the treatment effect sizes for short-run outcomes were sufficiently high that we were powered to detect reasonable long-run effects, we executed a follow-up survey. We recruited participants for the follow-up survey by sending a text message informing them that they could earn an additional KSH 350 by completing another survey.

Before sending the text, participants were randomized on whether the recruitment text reminded participants about course content through a “priming” treatment. These primes were customized to the course to which the participant was exposed in the Main Experiment to remind participants about what they previously learned. Figure 3 shows an example of the text message a participant in the Emotions course received with the prime. Online Appendix E contains the text for all primes.



🙄 Have you noticed any posts recently trying to create feelings of fear, anger, or superiority?! Just a quick reminder to STOP and QUESTION the information in the post when that happens. You don't want to share anything you're not 100% sure is true!

Want to earn 🇰🇸 KSH 350 🇰🇸 in mobile airtime TODAY? Take our 10-question survey now by replying START.

Figure 3: Message Sent to Emotions Course Participants Randomized into the Prime

2.2.1 Survey Structure

The follow-up survey showed participants five social media posts and asked them the same sharing question as in the pre- and post-survey. We did not ask the accuracy question (and so we did not implement the accuracy nudge) in the follow-up survey. Then, we asked five reflective questions:

1. What are some techniques that people use to create misleading social media posts?
2. When browsing your time line in the last month, did you notice any post that looked misleading? If so, what made it seem misleading?
3. How did you feel when you saw the misleading post? If you haven't seen any misleading posts recently, how do you think you would feel?
4. Has the Inoculation against Misinformation course changed your behavior on social media? If so, how?
5. If you were to tell a friend what you learned in the course, what tip would you share?

The five posts that participants saw were a non-misinformation post, an Emotions post, an Reasoning post, a Combo post, and an attention check post. Similar to how the post-survey matched misinformation posts to non-misinformation posts in the pre-survey, the follow-up matched misinformation posts to non-misinformation posts a participant saw in the post-survey.

2.3 Implementation

The Stage 1: Main Experiment was implemented in July 2022, with all participants completing the text message course intervention and post-survey by August 2022. The Stage 2: Follow-up Experiment was launched and concluded in September 2022. Table 2 shows the funnel of participants from starting the pre-survey to completing the follow-up survey. Attrition overall was fairly low due to the monetary incentive for completion, with 34% of participants who engaged with us completing the Main Experiment, and 61% of participants invited to the Follow-up Experiment (based on completing post-survey) completing it. More than 90% of participants who started each of our surveys (pre-, post-, or follow-up) completed them, suggesting the surveys were not too burdensome. Attrition during the text message course was higher, with only 66% of those who started the course completing it, but this attrition rate is still quite low considering we were implementing a five-day intervention.

	N Participants	% of "Started Pre-survey"	% of Previous Funnel Stage
Started Pre-survey	25287	-	-
Completed Pre-survey	22526	89.08%	89.08%
Started Text Message Course	18598	73.55%	82.56%
Completed Day 1 Course	16684	65.98%	89.71%
Completed Day 2 Course	13997	55.35%	83.89%
Completed Day 3 Course	13093	51.78%	93.54%
Completed Day 4 Course	11396	45.07%	87.04%
Completed Entire Course	10934	43.24%	95.95%
Started Post-survey	9589	37.92%	87.70%
Completed Post-survey	8678	34.32%	90.50%
Started Follow-up	5783	22.87%	66.64%
Completed Follow-up	5313	21.01%	91.87%

Table 2: Funnel of Participants

Notes: The number of participants who started the pre-survey is an upper bound estimate because we could count only survey copies on Qualtrics and not identify users. We were only able to identify users once they completed the pre-survey and provided a phone number, which is how we ensured that all down-funnel outcomes counted unique users based on phone number. The post- and follow-up surveys required users to validate their phone number before they could start the survey. We received 40,845 survey copies in total for the pre-survey, of which we discarded 3,092 users who had participated in one of our pilot studies and another 12,466 survey copies filled out by duplicated phone numbers. In the post-survey, we filtered out participants who encountered system errors (N = 104), did not have at least 5 days in between pre- and post-survey dates (N = 2,101), and/or did not complete the full text-message course (N = 509). For the follow-up survey, we filtered out 2,371 survey copies filled out by duplicated phone numbers and retained only the first copy done by each phone number.

Furthermore, Table 3 shows that attrition was not substantially different between assignment groups in the Main Experiment, with 36-43% of participants who completed the pre-survey completing the post-survey in all groups. Attrition was also not substantially different between assignment groups in the Follow-up Experiment, with 59-63% of participants who completed the post-survey also completing the follow-up survey in all groups. We show little difference in the proportion of participants who correctly answered attention check questions across intervention assignment groups; however, we do see sizeable differences between the Accuracy Inter and Accuracy After assignment groups. Participants in the Accuracy Inter group are consistently six to nine percentage points more likely to correctly answer both attention check questions, suggesting that the Accuracy

Inter participants were paying closer attention. In Appendix D.2, we show that our results are robust to restricting our sample to those who correctly answer the attention check questions.

	No-course baseline	Facts baseline	Reasoning course	Emotions course	Combo course	Totals
Phase 1: Main Exp.						
Accuracy After	825 (36.38%) [31.03%]	885 (39%) [32.99%]	856 (37.93%) [31.89%]	894 (38.67%) [33.45%]	858 (37.78%) [34.97%]	4318 (37.95%) [32.89%]
Accuracy Inter	800 (36%) [37.38%]	904 (39.95%) [41.26%]	834 (37.72%) [40.89%]	942 (42.72%) [40.76%]	880 (39.15%) [40.68%]	4360 (39.11%) [40.25%]
Totals	1625 (36.19%) [34.15%]	1789 (39.47%) [37.17%]	1690 (37.82%) [36.33%]	1836 (40.65%) [37.2%]	1738 (38.46%) [37.86%]	8678 (38.52%) [36.59%]
Phase 2: Follow-up Exp.						
	962 (59.16%) [85.45%] {43.24%}	1069 (59.75%) [86.25%] {44.53%}	1072 (63.43%) [86.85%] {41.14%}	1122 (61.11%) [84.05%] {42.6%}	1088 (62.6%) [85.85%] {43.84%}	5313 (61.24%) [85.68%] {43.09%}

Table 3: Attrition and Attention Check Passing Rates, by Intervention Assignment Group

Notes: Main cells show the number of participants who completed the post-survey (in the Main Experiment) and the number of participants who completed the follow-up (in the Follow-up Experiment) in their respective assignment group. Numbers in parentheses are the ratios of the number of participants who completed the post-survey to the number of participants who completed the pre-survey (in the Main Experiment) and the ratios of the number of participants who completed the follow-up survey to the number of participants who completed the post-survey (in the Follow-up Experiment). Numbers in brackets are the ratios of the number of participants who correctly answered all attention checks in the pre- and post-survey to the number of participants who completed the post-survey (in the Main Experiment) and the ratios of the number of participants who correctly answered the single attention check in the follow-up survey to the number of participants who completed the follow-up survey (in the Follow-up Experiment). Numbers in curly brackets are the ratios of the number of participants who correctly answered all attention check questions in the pre-, post-, and follow-up surveys to the number of participants who completed the follow-up survey.

The descriptive statistics in Table 4 show that our sample is young at approximately 26 years old, disproportionately men, mostly single, and predominantly Christian. Participants are more likely to be unemployed than employed (part-time, full-time, or self-employed), and they are about evenly split between rural, urban, and suburban communities.

	Mean	Std. Dev	Q1	Median	Q3	% Missing
Age	26.40	7.593	22	25	29	0%
Man	65.42%	0.476				0%
Educational Attainment						
High school or less	27.1%	0.444				0%
Some college	35.56%	0.479				0%
Bachelor's degree	36.07%	0.48				0%
Graduate degree	1.27%	0.112				0%
Married	32.57%	0.469				0%
Employment Status						
Unemployed	40.09%	0.49				0%
Employed	32.52%	0.468				0%
Student	27.39%	0.446				0%
Location						
Mostly urban	28.51%	0.451				0%
Suburban	38.53%	0.487				0%
Mostly rural	32.96%	0.47				0%
Christian	94.19%	0.234				0%
Attends religious services	94.63%	0.225				0%
Uses social media	99.33%	0.081				0%
Hrs/day on social media	5.82	3.696	3	5	7	0.67%
Prop. of content shared						
0-20%	16.94%	0.375				0%
20-40%	20%	0.4				0%
40-60%	30.63%	0.461				0%
60-80%	23.3%	0.423				0%
80-100%	8.45%	0.278				0%
Pre-survey Sharing						
All posts	0.62	0.326	0.444	0.667	0.889	0%
Non-misinformation posts	0.68	0.362	0.333	0.667	1	0%
Misinformation posts	0.59	0.347	0.333	0.667	1	0%

Table 4: Descriptive Statistics for Post-Survey Completers Sample

Notes: Sample includes 8,768 participants who completed the post-survey. See Appendix B3 for definitions of covariates.

Appendix B.3 contains figures that show attrition and balance. We first consider what kind of selection on observables we observe based on attrition in the post- and follow-up surveys. We observe no differences larger than a tenth of a standard deviation between the participants who completed the post-survey and those who completed the pre-survey only. We observe a few differences between those who completed the follow-up survey and those who completed the post-survey only. Follow-up completers are younger, less-educated, and more likely to be unmarried and to be a student. Then, we consider the balance across the text message course intervention assignment groups. Figure B2 shows that there were no differences greater than a tenth of a standard deviation between assignment groups among participants who completed the post-survey. Importantly, we observed no differences in sharing behavior in the pre-survey.

We do, however, see two small differences among participants who complete the follow-up survey. Participants in the Reasoning course group are less likely to be married and more likely to share non-misinformation posts relative to participants in the Facts baseline. The latter difference could be concerning as it determines the denominator in the primary outcome of interest; therefore, we

include Augmented Inverse Propensity Weighted (AIPW) estimates for our results in Appendix D and find that they are robust to controlling for observed covariates.

3 Results

This section evaluates the results of the experiment. Section 3.1 defines the outcomes of interest and presents summary statistics for outcomes by treatment assignment status. Section 3.2 analyzes treatment effects for the text message courses. Section 3.3 presents results with alternative outcome measures, while Section 3.4 establishes that the results persist in the long run. Section 3.5 analyzes the accuracy nudge. In Appendix A, we present all of the results specified in our pre-analysis plan (Athey et al., 2022).

3.1 Defining and Summarizing Outcomes

To define the primary outcome, let $S_i^{N,pre}(j)$ be participant i 's sharing decision for the non-misinformation N post about fact $j \in \{1, 2, 3\}$ in the pre-survey and $S_i^{M,post}(j)$ their sharing decision for the misinformation M post about fact j in the post-survey. Then, our sharing outcome is constructed as follows:

$$share_i = \frac{\sum_{j=1}^3 1 \left[S_i^{M,post}(j) | S_i^{N,pre}(j) = 1 \right]}{\sum_{j=1}^3 S_i^{N,pre}(j)} \quad (1)$$

The denominator is the number of non-misinformation posts participant i shares in the pre-survey out of the three they are shown. The numerator is the number of misinformation posts corresponding to the non-misinformation posts participant i shares in the pre-survey that participant i shares in the post-survey (recall the example in Table 1). So, a participant who shares all three non-misinformation posts in the pre-survey can share $\{0, 1, 2, 3\}$ corresponding posts in the post-survey and their outcome can take on the value $\{0, \frac{1}{3}, \frac{2}{3}, 1\}$. A participant who shares only two of the non-misinformation posts in the pre-survey can share $\{0, 1, 2\}$ corresponding posts in the post-survey, where the observation of the misinformation post corresponding to the non-misinformation post that the participant does not share in the pre-survey is dropped, and their outcome can take on the value $\{0, \frac{1}{2}, 1\}$. 994 participants who share zero of the non-misinformation posts in the pre-survey are dropped from the analysis for this outcome. Because sharing the pre-survey differs between accuracy nudge groups, we differentially drop participants in the Accuracy Inter (656) and Accuracy After groups (338). This differential dropping is not problematic because accuracy nudge groups are randomized into text message course assignment groups.

This outcome is designed to disentangle the part of the sharing decision that is related to the participant's evaluation of a post as misinformation from the part that is related to the participant's

interest in the post’s content. By measuring baseline preferences for post content with no misinformation, we can isolate the part of the sharing decision that is based on the participant’s evaluation of a post as misinformation alone. Subsection 3.3 shows that participants are 43.8% less likely to share misinformation posts corresponding to non-misinformation posts they do *not* intend to share in the pre-survey than misinformation posts corresponding to non-misinformation posts they do intend to share (22.5% compared to 51.5%, on average), indicating that participant interest in content is an important driver of variation in sharing overall. Thus, accounting for interest can be helpful for reducing variation in our outcome measures.

In addition to our primary outcome, we also consider outcomes that relate to the sharing of posts that are not misinformation. If our treatments affect sharing of all posts equally, then they cannot be interpreted as specifically targeting misinformation. An outcome commonly used in the literature is sharing *discernment*, which is the difference between the number of non-misinformation posts shared and the number of misinformation posts shared.¹⁸ Another discernment outcome we consider takes the difference in the accuracy score of non-misinformation posts and the accuracy score of the misinformation posts. Section 3.3 examines these and additional outcome measures that do not include the posts used to construct the primary outcome

Table 5 presents means and standard errors for a variety of outcomes, broken out by both the text message course assignment group and assignment to the accuracy nudge. Outcomes include Sharing Rate (Equation 1), discernment, and the proportion of misinformation and non-misinformation posts shared in the pre-survey and the post-survey. Table 5 forms the basis of the treatment effect estimates presented in subsequent sections.

3.2 Text Message Treatment Course Effectiveness

We answer the following two questions to evaluate the text message treatment courses’ effectiveness. Do the treatment course have different outcomes than our baselines? And do outcomes differ across the Emotions, Reasoning, and Combo treatment courses? Figure 4 illustrates averages for the Sharing Rate outcome across each of the text message course interventions for the 7,774 participants who both completed the post-survey and shared at least one non-misinformation post in the pre-survey.

Several issues inform the interpretation of the results. First, note that our sample is selected based on completing the post-survey, and so treatment effects should be interpreted in that context. However, as described in Section 2.3, selected individuals have similar observable covariates to the overall recruited population, and are comparable across treatment arms.

Second, as we prespecified, other than when we analyze the accuracy nudge specifically in subsection 3.5, we pool the two accuracy nudge groups when analyzing the text message course. The detailed breakdowns for each accuracy nudge condition are included in Table 5. When interpreting pooled results for the Sharing Rate, note that Table 5 shows that the accuracy questions interweaved between the sharing questions in the Accuracy Inter group decrease misinformation sharing relative

¹⁸Offer-Westort et al. (2021) studies a variant where true posts are given weight .5 relative to false posts receiving weight 1.

	No-course baseline	Facts baseline	Reasoning course	Emotions course	Combo course	Treatment courses	Baselines	All
Number of Observations								
Accuracy Inter	826	886	856	894	859	2609	1712	4321
Accuracy After	800	903	833	942	879	2654	1703	4357
All	1626	1789	1689	1836	1738	5263	3415	8678
Pre Non-misinfo Posts								
Accuracy Inter	0.651 (0.013)	0.628 (0.012)	0.655 (0.012)	0.642 (0.012)	0.647 (0.012)	0.648 (0.007)	0.639 (0.009)	0.644 (0.005)
Accuracy After	0.756 (0.01)	0.739 (0.011)	0.752 (0.011)	0.738 (0.011)	0.726 (0.011)	0.739 (0.006)	0.747 (0.007)	0.742 (0.005)
All	0.704 (0.008)	0.683 (0.008)	0.704 (0.008)	0.689 (0.008)	0.686 (0.008)	0.693 (0.005)	0.693 (0.006)	0.693 (0.004)
Pre Misinfo Posts								
Accuracy Inter	0.525 (0.012)	0.498 (0.011)	0.503 (0.012)	0.519 (0.011)	0.518 (0.012)	0.514 (0.007)	0.511 (0.008)	0.512 (0.005)
Accuracy After	0.605 (0.01)	0.581 (0.01)	0.576 (0.011)	0.563 (0.011)	0.574 (0.011)	0.571 (0.006)	0.592 (0.007)	0.58 (0.005)
All	0.566 (0.008)	0.539 (0.008)	0.54 (0.008)	0.541 (0.008)	0.546 (0.008)	0.542 (0.005)	0.552 (0.006)	0.546 (0.004)
Post Non-misinfo Posts								
Accuracy Inter	0.639 (0.013)	0.559 (0.013)	0.518 (0.013)	0.477 (0.012)	0.486 (0.013)	0.493 (0.007)	0.596 (0.009)	0.533 (0.006)
Accuracy After	0.719 (0.012)	0.657 (0.012)	0.582 (0.013)	0.563 (0.012)	0.559 (0.013)	0.568 (0.007)	0.687 (0.008)	0.615 (0.006)
All	0.68 (0.009)	0.607 (0.009)	0.551 (0.009)	0.519 (0.009)	0.522 (0.009)	0.53 (0.005)	0.642 (0.006)	0.574 (0.004)
Post Misinfo Posts								
Accuracy Inter	0.514 (0.012)	0.428 (0.012)	0.38 (0.012)	0.326 (0.01)	0.356 (0.012)	0.353 (0.006)	0.468 (0.009)	0.398 (0.005)
Accuracy After	0.578 (0.011)	0.504 (0.012)	0.427 (0.011)	0.381 (0.011)	0.387 (0.011)	0.398 (0.006)	0.539 (0.008)	0.454 (0.005)
All	0.546 (0.008)	0.465 (0.008)	0.404 (0.008)	0.353 (0.007)	0.372 (0.008)	0.376 (0.005)	0.504 (0.006)	0.426 (0.004)
Primary Outcome: Sharing Rate								
Accuracy Inter	0.63 (0.015)	0.544 (0.015)	0.495 (0.015)	0.413 (0.014)	0.462 (0.015)	0.455 (0.009)	0.584 (0.011)	0.505 (0.007)
Accuracy After	0.647 (0.013)	0.582 (0.014)	0.498 (0.014)	0.447 (0.014)	0.447 (0.014)	0.464 (0.008)	0.614 (0.01)	0.524 (0.006)
All	0.639 (0.01)	0.563 (0.01)	0.496 (0.01)	0.43 (0.01)	0.454 (0.01)	0.459 (0.006)	0.6 (0.007)	0.515 (0.005)
Pre Sharing Discernment								
Accuracy Inter	-1.199 (0.058)	-1.101 (0.053)	-1.052 (0.057)	-1.188 (0.053)	-1.167 (0.052)	-1.138 (0.031)	-1.147 (0.039)	-1.142 (0.024)
Accuracy After	-1.361 (0.053)	-1.269 (0.051)	-1.201 (0.053)	-1.167 (0.055)	-1.268 (0.055)	-1.211 (0.031)	-1.313 (0.037)	-1.252 (0.024)
All	-1.281 (0.039)	-1.184 (0.037)	-1.127 (0.039)	-1.178 (0.038)	-1.217 (0.038)	-1.174 (0.022)	-1.23 (0.027)	-1.196 (0.017)
Post Sharing Discernment								
Accuracy Inter	-1.165 (0.056)	-0.89 (0.053)	-0.727 (0.053)	-0.525 (0.047)	-0.681 (0.052)	-0.641 (0.029)	-1.019 (0.038)	-0.789 (0.023)
Accuracy After	-1.31 (0.053)	-1.05 (0.053)	-0.818 (0.056)	-0.597 (0.052)	-0.646 (0.053)	-0.686 (0.031)	-1.175 (0.038)	-0.88 (0.024)
All	-1.239 (0.038)	-0.969 (0.038)	-0.773 (0.038)	-0.56 (0.035)	-0.664 (0.037)	-0.663 (0.021)	-1.098 (0.027)	-0.834 (0.017)
Pre Accuracy Score Discernment								
Accuracy Inter	1.91 (0.255)	2.395 (0.233)	2.813 (0.249)	2.176 (0.227)	2.283 (0.238)	2.411 (0.137)	2.167 (0.172)	2.316 (0.107)
Accuracy After	1.215 (0.24)	1.445 (0.223)	1.451 (0.231)	1.843 (0.235)	1.799 (0.237)	1.7 (0.135)	1.334 (0.163)	1.555 (0.104)
All	1.557 (0.175)	1.925 (0.162)	2.123 (0.171)	2.014 (0.163)	2.044 (0.168)	2.059 (0.097)	1.75 (0.119)	1.937 (0.075)
Post Accuracy Score Discernment								
Accuracy Inter	1.772 (0.236)	3.317 (0.233)	3.651 (0.229)	4.142 (0.209)	3.853 (0.227)	3.892 (0.128)	2.591 (0.167)	3.384 (0.102)
Accuracy After	1.051 (0.234)	2.28 (0.223)	2.916 (0.227)	3.347 (0.214)	3.165 (0.223)	3.146 (0.128)	1.687 (0.162)	2.568 (0.101)
All	1.406 (0.166)	2.803 (0.162)	3.278 (0.161)	3.755 (0.15)	3.513 (0.159)	3.522 (0.091)	2.138 (0.117)	2.977 (0.072)

Table 5: Summary of Outcomes by Text Message Course Intervention and Accuracy Nudge Assignment

Notes: Sample includes 8,768 participants who completed the post-survey, except for the “Primary Outcome: Sharing Rate” rows that exclude 994 participants who did not share at least one non-misinformation post in the pre-survey. The first two rows display the number of observations in each assignment group. The other rows display averages by assignment group, with standard errors in parentheses.

to the Accuracy After group in both the pre- and post-surveys.¹⁹ This result means that the distribution of our denominator that serves as a baseline in our primary outcome (as well as the numerator) differ between the two accuracy nudge groups. Since assignment to the two accuracy conditions is balanced across the text message course assignments, this issue does not bias our results.

Figure 4 illustrates contrasts between the treatment courses and baselines. Relative to the No-course baseline sharing rate of 63.9 p.p., the treatment courses decrease the Sharing Rate by 18.1 p.p. ($SE = 1.16$ p.p.) or approximately 28%. Relative to the Facts baseline, the treatment course decreases the Sharing Rate by 10.4 p.p. ($SE = 1.19$ p.p.) or approximately 18% relative to the 56.3 p.p. sharing rate for the Facts baseline. The difference between the No-course and Facts baselines of 7.7 p.p. ($SE = 1.44$ p.p.) shows that making misinformation salient to participants on a daily basis does have some beneficial impact on sharing behavior, but the treatment courses are providing educational content beyond this salience effect.




Figure 4: Misinformation Sharing, by Intervention Assignment Group

Notes: Sample includes 7,774 participants who completed the post-survey and shared at least one non-misinformation post in the pre-survey. Each bar displays the Sharing Rate for misinformation posts, as defined in Equation 1, by participants in their respective intervention assignment group, pooling participants in the Accuracy Inter and Accuracy After groups. Above each bar, the standard error is shown in parentheses below the Sharing Rate. The thin black bars represent 95% confidence intervals. Differences in Sharing Rates are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

Figure 5 shows that the best performing course is the Emotions course. Further, the Emotions course reduces misinformation sharing by 2.3 p.p. ($SE = 1.44$) more than the Combo course. Thus, we can bound any benefit of the Combo course to no more than 0.5 p.p., which is 1.1% of the 45.8 p.p. baseline sharing rate in the treatment courses in the post-survey or 2.8% of the 18.1 p.p. treatment

¹⁹We discuss the estimates for the differences between the accuracy nudge groups in subsection 3.5.

effect of the treatment courses relative to the No-course baseline. This result implies that teaching the reasoning-based techniques in the Combo course is, at best, not harmful.²⁰ In comparison to the Reasoning course, the Emotions course decreases misinformation sharing by 6.6 p.p. ($SE = 1.43$ p.p.) more. The Combo course also decreases misinformation sharing more than the Reasoning course by 4.3 p.p. ($SE = 1.47$), further confirming the importance of focusing on emotion-based techniques. In summary, teaching emotion-based techniques decreases misinformation sharing more than teaching reasoning-based techniques, and the reasoning-based content has no added value that we can detect when combined with the emotion-based content.




Figure 5: Misinformation Sharing for Treatment Courses Only, by Intervention Assignment Group

Notes: Sample includes 4,646 participants in the treatment courses who completed the post-survey and shared at least one non-misinformation post in the pre-survey. Each bar displays the Sharing Rate for misinformation posts, as defined in Equation 1, by participants in their respective intervention assignment group, pooling participants in the Accuracy Inter and Accuracy After groups. Above each bar, the standard error is shown in parentheses below the Sharing Rate. The thin black bars represent 95% confidence intervals. Differences in Sharing Rates are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

3.2.1 Heterogeneity by Misinformation Post Type

This section explores the effectiveness of the different treatment courses for different types of posts. In the pre- and post-surveys, we show participants multiple types of misinformation, defined in the same way as the treatment courses: Emotions, Reasoning, and Combo. This correspondence allows us to evaluate whether, for example, participants in the Emotions course are less likely to intend to share Emotions posts, but are no different with respect to Reasoning posts. In particular, based on

²⁰Note that when we report p-values for hypothesis tests in Appendix A, the hypotheses tested in this paragraph comparing the treatment courses to one another are grouped in Family 2 with tests broken out by the type of post presented in Figure 6 when correcting for multiple hypothesis testing.

the result from section 3.2 that the Emotions course outperforms the Reasoning course, we want to learn whether the Reasoning course might outperform the Emotions course in Reasoning posts in particular, even if it performs worse overall. Understanding the cross-technique effect of the treatment courses also contributes to our understanding of how the strategies taught in the courses generalize to techniques not taught.

Figure 6 shows that, not only does the Emotions course outperform the Reasoning and Combo courses in Emotions posts, it does not perform statistically worse on Reasoning or Combo posts. The Emotions course decreases the sharing of Emotions posts by 7.9 p.p. ($SE = 2.02$ p.p.) more than the Reasoning course and by 5.7 p.p. ($SE = 2.02$ p.p.) more than the Combo course. Furthermore, the Emotions course decreases misinformation sharing on Combo posts by 5.0 p.p. ($SE = 2.91$) more than the Combo course and by 7.7 p.p. ($SE = 1.99$ p.p.) more than the Reasoning course. While the Emotions course also decreases misinformation on Reasoning posts by 3.0 p.p. ($SE = 2.02$ p.p.) more than the Reasoning course, the Combo course outperforms the Emotions course on Reasoning posts by 1.9 p.p. ($SE = 2.01$ p.p.). We can rule out that the Combo course outperforms the Emotions course by more than 5.9 p.p. on Reasoning posts, and the difference is not statistically distinguishable even with our large sample size ($p = 0.452$). Appendix Table A3 shows the differences for each prespecified pairwise comparison and the Romano-Wolf adjusted p-values. We fail to reject differences between the Emotions course and the Combo course on both the Reasoning and Combo posts as well as the difference between the Emotions course and the Reasoning course on Combo posts.

The result that the Emotions course is best on Emotions posts, and not detectably worse than the other two treatment courses on Reasoning and Combo posts, confirms that the Emotions course is the most effective of the interventions we test. The strategy targeting the first-stage of the decision-making process is changing sharing behavior on all types of misinformation we test, not just the misinformation posts that use the emotion-based techniques taught in the Emotions course.

3.3 Alternative Outcome Measures

We test four alternative outcome measures in this section to assess the robustness of our results. The first three alternative outcomes vary how we measure misinformation sharing. First, we test one of the standard outcome measures used in the literature: sharing discernment. Second, we evaluate misinformation sharing using only misinformation posts not used in the main analysis. Third, we evaluate the “opposite” of our primary outcome in Equation 1 by sub-setting to non-misinformation posts that participants do NOT intend to share in the pre-survey (instead of sub-setting to the posts participants intend to share). The fourth alternative outcome uses the accuracy scores we elicit in addition to sharing decisions.

The first alternative outcome we assess is sharing discernment, which is computed by taking the difference between the number of non-misinformation posts and the number of misinformation posts a participant intends to share. This outcome is commonly used in the literature to handle concerns that interventions to decrease misinformation sharing may reduce sharing overall, not just sharing of misinformation. Positive values mean the participant is sharing more non-misinformation than misinformation posts, so taking the post- minus pre-survey difference means that more positive




Figure 6: Misinformation Sharing, by Intervention Assignment Group and Post Type

Notes: Sample includes 4,646 participants in the treatment courses who completed the post-survey and shared at least one non-misinformation post in the pre-survey. Each group of bars displays the Sharing Rate for misinformation posts of each type, as defined in Equation 1, by participants in their respective intervention assignment group, pooling participants in the Accuracy Inter and Accuracy After groups. Above each bar, the standard error is shown in parentheses below the Sharing Rate. The thin black bars represent 95% confidence intervals. Differences in Sharing Rates are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

values represent *better* discernment.

Figure 7 shows that all treatment courses improve discernment relative to both the No-course baseline and the Facts baseline, with Emotions performing the best of the three treatment courses. Relative to a baseline discernment of -1.28 ($SE = 0.055$), the Emotions course increases discernment by 0.40 more than the Facts baseline. As expected, participants in the No-course baseline show no change in discernment.

The second alternative outcome is the proportion of misinformation posts shared, excluding the three misinformation posts in the post-survey that correspond to non-misinformation posts in the pre-survey. In the pre-survey, the denominator for this measure is six; while for the post-survey, the denominator is only three because three of the six misinformation posts participants saw in the post-survey are used in the primary outcome. The concern we address with this outcome is that seeing the same fact twice (as non-misinformation, then as misinformation) could affect our results through memory, although we posit that the five days in between negate that concern. By using only misinformation posts for which participants saw the related fact just once, we can assess the robustness of our results to this concern. We compare the post- minus pre-survey difference, so negative values mean sharing *less* misinformation in the post-survey.




Figure 7: Differences in Discernment Score, by Intervention Assignment Group

Notes: Sample includes 8,768 participants who completed the post-survey. Each point displays the average post-minus pre-survey difference in sharing discernment, defined as the number of non-misinformation posts shared minus the number of misinformation posts shared, for participants in their respective intervention assignment group, pooling participants in the Accuracy After and Accuracy Inter groups. Above each point, the standard error of each difference is shown in parentheses below the difference. The thin colored bars represent 95% confidence intervals. Differences in the average differences are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

Figure 8 shows that, again, all treatment courses do better than both baselines at decreasing misinformation sharing. The Emotions course has the largest treatment effect, but it is only 1.6 p.p. ($SE = 1.26$ p.p.) larger than the treatment effect for the Combo course. Note that there is some decrease in misinformation sharing in the No-course baseline of 2.7 p.p. ($SE = 0.80$ p.p.), indicating a small change in behavior between the pre- and post-survey even without exposure to a text message course. One possible explanation is that the accuracy questions at the end of the pre-survey affect sharing behavior in the post-survey. The decrease in misinformation sharing from the pre- to the post-survey in the Accuracy After group is 3.8 p.p. ($SE = 1.23$ p.p.) compared to 1.5 p.p. ($SE = 1.32$ p.p.) in the Accuracy Inter group, which is consistent with this explanation.

We define our third alternative outcome, which we call the opposite outcome because its denominator is the complement of the denominator for the primary outcome in Equation 1, as follows:

$$share_i = \frac{\sum_{j=1}^3 1 \left[S_i^{M,post}(j) | S_i^{T,pre}(j) = 0 \right]}{3 - \sum_{l=1}^3 S_i^{T,pre}(l)} \quad (2)$$

Here the denominator is the number of non-misinformation posts participant i does NOT intend




Figure 8: Differences in Non-Corresponding Misinformation Sharing, by Intervention Assignment Group

Notes: Sample includes 8,768 participants who completed the post-survey. Each point displays the average post-minus pre-survey difference in the proportion of misinformation posts shared, excluding those misinformation posts in the post-survey that correspond to non-misinformation posts in the pre-survey, for participants in their respective intervention assignment group, pooling participants in the Accuracy After and Accuracy Inter groups. Above each point, the standard error of each difference is shown in parentheses below the difference. The thin colored bars represent 95% confidence intervals. Differences in the average differences are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

to share in the pre-survey out of the three they are shown. The numerator is the number of misinformation posts corresponding to the non-misinformation posts participant i does NOT intend to share in the pre-survey that participant i intends to share in the post-survey. We drop 4,069 participants who intend to share all non-misinformation posts because then the denominator of this outcome equals zero and is undefined. Because sharing behavior differs between the accuracy nudge groups, we drop more Accuracy After participants than Accuracy Inter participants. Again, this differential dropping is not problematic because accuracy nudge groups are randomized into text message course assignment groups.

This outcome is motivated by the concern that people may be more likely to share misinformation than non-misinformation, holding constant the relevant fact, *and* that the treatment courses might also differentially affect sharing based on interest in the content versus sharing based on a preference for misinformation over non-misinformation posts. Figure 9 shows that all treatment courses are an improvement over the No-course and Facts baselines using this outcome measure, but the treatment effects are much smaller. This decrease in magnitude is partially due to lower rates in baselines (63.9% on average for the primary outcome versus 29.5% for this outcome), but the treatment effect size in percentage term is also lower. The primary outcome decreases 24% from the Facts baseline to the Emotions course, but the opposite outcome decreases only 17% in the same comparison.

This result suggests that it may be more difficult to have an impact with the treatment courses on people who are more likely to share a misinformation version of a post than a non-misinformation version of the same post.




Figure 9: Opposite Outcome, by Intervention Assignment Group

Notes: Sample includes 4,699 participants who completed the post-survey and shared at most two non-misinformation posts in the pre-survey. Each bar displays the Opposite Sharing Rate for misinformation posts, as defined in Equation 2, by participants in their respective intervention assignment group, pooling participants in the Accuracy Inter and Accuracy After groups. Above each bar, the standard error is shown in parentheses below the Opposite Sharing Rate. The thin black bars represent 95% confidence intervals. Differences in Opposite Sharing Rates are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

The baseline for this alternative outcome is also a validation of our primary outcome. The implicit assumption in our primary outcome is that, if a participant does not intend to share a fact in a non-misinformation post, they will be less likely to report an intention to share that same fact in a misinformation post as well. If we instead observe that misinformation sharing rates are higher on facts that participants do not intend to share in the pre-survey compared to those that participants do intend to share, we would be concerned about the validity of our primary outcome. Comparing Figure 4 to Figure 9 shows that, as expected, participants are more likely to indicate an intention to share misinformation posts about a fact for which they have already indicated an intention to share a non-misinformation post (51.5% on average) compared to facts for which they have already indicated an intention to NOT share a non-misinformation post (22.5% on average), validating the assumption underlying our primary outcome.

All outcomes and tests on sharing behavior support the conclusion that the treatment courses, and the Emotions course in particular, effectively decrease misinformation sharing; however, we might be concerned that these interventions are only stopping people from sharing posts that could be misinformation, but not changing how they think about the misinformation posts. We address this

concern with the last alternative outcome using the accuracy scores participants give for the same posts for which they report their sharing decision. Accuracy scores are defined using a numerical encoding of the answer to the “To the best of your knowledge, how accurate is the claim in the above post?” where “Not at all accurate” is encoded as -3 or 3 if the post is a non-misinformation post or a misinformation post, respectively; “Not very accurate” is encoded as -1 or 1 if the post is a non-misinformation post or a misinformation post, respectively; “Somewhat accurate” is encoded as 1 or -1 if the post is a non-misinformation post or a misinformation post, respectively; and “Very accurate” is encoded as 3 or -3 if the post is a non-misinformation post or a misinformation post, respectively. Summing this numerical encoding of all posts a participant saw in either the pre- or the post-survey yields yields accuracy discernment.




Figure 10: Accuracy Score Discernment, by Intervention Assignment Group

Notes: Sample includes 8,768 participants who completed the post-survey. Each point displays the average post-minus pre-survey difference in accuracy discernment for participants in their respective intervention assignment group, pooling participants in the Accuracy After and Accuracy Inter groups. Above each point, the standard error of each sharing rate is shown in parentheses below the sharing rate. The thin colored bars represent 95% confidence intervals. Differences in the average sharing rates are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses below.

Figure 10 shows that the treatment courses increase accuracy discernment over both baselines. The Emotions course again has the largest treatment effect, increasing discernment by 0.86 ($SE = 0.279$) compared to 0.59 ($SE = 0.287$) in the Combo course and 0.27 ($SE = 0.290$) in the Reasoning course. Accuracy discernment in the No-course baseline does not change, but making misinformation salient for five days improves accuracy discernment by 1.03 ($SE = 0.294$) in the Facts baseline. In summary, the treatment courses not only decrease misinformation sharing behavior, they decrease how accurate participants perceive misinformation to be.

3.4 Long-run effects

We test the long-run effects of the treatment courses in this section. To evaluate the persistent effects of the treatment courses, we again use our primary outcome defined in Equation 1, but instead use the non-misinformation posts in the post-survey and their corresponding misinformation posts in the follow-up survey. After dropping participants who did not share any non-misinformation post in the follow-up survey, our sample includes 4,249 participants who completed the follow-up survey. Note that again, we have differential dropping between accuracy nudge groups, but these groups are randomized into the treatment courses we evaluate.

Figure 11 shows that all treatment courses have a persistent effect on misinformation sharing seven to eleven weeks after course completion. The Emotions course continues to be the most effective, decreasing misinformation sharing by 12.3 p.p. ($SE = 1.98$ p.p) more than the Facts baseline. In comparison, the Reasoning and Combo courses decrease misinformation sharing by only 6.8 p.p. ($SE = 1.98$ p.p.) and 8.1 p.p. ($SE = 2.00$ p.p.) more than the Facts baseline, respectively.




Figure 11: Misinformation Sharing in Follow-up, by Intervention Assignment Group

Notes: Sample includes 4,249 participants who completed the follow-up survey and shared at least one non-misinformation post in the post-survey. Each bar displays the Sharing Rate for misinformation posts, as defined in Equation 1, changing the *post* and *pre* superscripts to *followup* and *post*, respectively, by participants in their respective intervention assignment group, pooling participants in the Accuracy Inter and Accuracy After groups. Above each bar, the standard error is shown in parentheses below the Sharing Rate. The thin black bars represent 95% confidence intervals. Differences in the average Sharing Rates are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

Recall that the No-course baseline group received the Combo course after completing the post-survey, which raises the question of why the Combo course assignment group decreases misinformation sharing 8.18 p.p. ($SE = 2$ p.p.) more than the No-course baseline even though both are

exposed to the same intervention at the time of the Follow-up Experiment. This difference is likely due to low take-up (52%) of the Combo course because payment was already complete. Given that participants in other text message course intervention groups had to complete the text message course prior to completing the post-survey, 100% of the sample in the Combo course assignment group was exposed to the text message course. In comparison, only about half of the sample in the No-course baseline was exposed to the text message course.

We pool across participants assigned to the prime in Figure 11 because the difference in sharing rates between primed and non-primed participants is only 2.0 p.p. ($SE = 1.82$ p.p.). The persistent effect of the treatment courses in combination with a null effect of the prime suggests that the treatment courses are effectively educating users. Appendix Figures B6 and B7 show the results for the follow-up separately by primed and non-primed participants.

Figure 12 shows that we continue to find that the Emotions course decreases misinformation sharing on Emotions posts more than the other courses (7.95 p.p., $SE = 2.90$ p.p., more than the Reasoning course, and 6.93 p.p., $SE = 2.90$ p.p., more than the Combo course). The differences between the Emotions course and the other courses on Reasoning and Combo posts have the same estimated direction as in the Main Experiment, except for the difference with the Combo course on the Reasoning posts, as Emotions decreases misinformation sharing by 2.87 ($SE = 2.96$ p.p.) more than the Combo course in the follow-up survey.

3.4.1 Reflective Questions

Last, we analyze the responses to two of the reflective questions participants answer in the follow-up survey: “Has the Inoculation against Misinformation course changed your behavior on social media? If so, how?” (Question 4 or Q4) and “If you were to tell a friend what you learned in the course, what tip would you share?” (Question 5 or Q5). The follow-up survey includes these questions so that we could elicit from participants what they report learning from the course and integrating into their social media behavior. As discussed in the Introduction, arguably the most important limitation of this study is its use of survey sharing outcomes, so we use these free response questions to learn about behavior outside of the survey. This approach has limitations, as we only know what participants are reporting, which could still be subject to experimenter demand.

To start, Table 6 provides summary statistics on all five questions we include in our analysis. Responses are 65.6 characters ($SE = 0.63$), or 12.2 words ($SE = 0.11$), on average for Q4 and 68.1 characters ($SE = 0.59$), or 11.6 words ($SE = 0.10$), for Q5, both of which are longer responses than the other three questions. In general, however, participants wrote a meaningful quantity of text. With the notable exception of in Question 3 asking about participants’ feelings when they see misinformation, participants in the Emotions course wrote the most. On the two questions we include in this analysis, participants in the Emotions course wrote 3.3 characters ($SE = 1.82$) and 2.9 characters ($SE = 1.91$) more compared to the Facts baseline.

To learn whether participants in treatment courses are more likely to report behavior changes on social media consistent with treatment, we use keywords like “stop,” “pause,” “question,” and “evaluate” to define a binary variable based on whether a participant’s response to the question

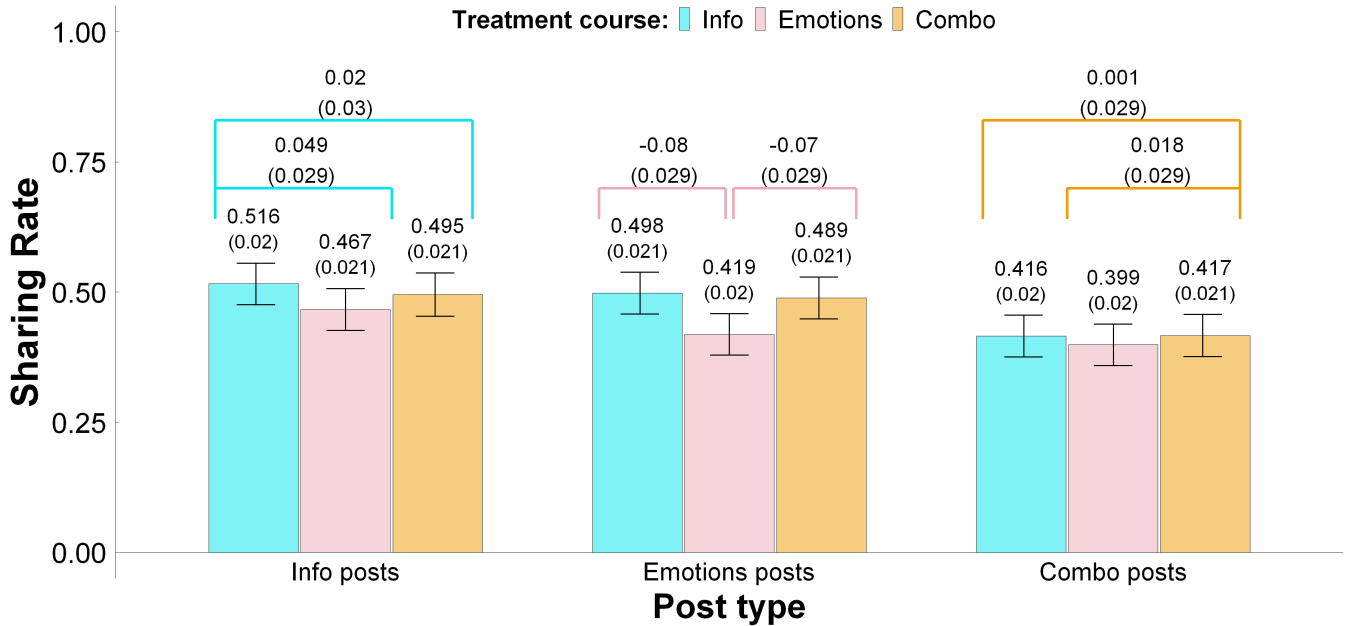


Figure 12: Misinformation Sharing in Follow-up, by Intervention Assignment Group and Post Type

Notes: Sample includes 2,555 participants in the treatment courses who completed the follow-up survey and shared at least one non-misinformation post in the post-survey. Each group of bars displays the Sharing Rate for misinformation posts of each type, as defined in Equation 1, changing the *post* and *pre* superscripts to *followup* and *post*, respectively-, by participants in their respective intervention assignment group, pooling participants in the Accuracy Inter and Accuracy After treatments. Above each bar, the standard error is shown in parentheses below the Sharing Rate. The thin black bars represent 95% confidence intervals. Differences in the average Sharing Rates are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

includes one of these keywords. These words correspond to discussing a technique taught in one or more of the courses. Figure 13 shows that these keywords are more likely to show up in responses from participants in one of the treatment courses compared to the Facts baseline.²¹ Figure 14 shows that participants in the treatment courses are also more likely to tell their friends a tip with these keywords than participants in the Facts baseline. In summary, participants in the treatment courses are more likely to report behavior consistent with the strategies suggested in the treatment courses than the Facts baseline.

3.5 Accuracy Nudge

In this section, we address three questions. First, does the accuracy nudge effectively decrease misinformation sharing in our study? Second, is the accuracy nudge more or less effective than

²¹It is not surprising that these words also come up in responses from participants in the Facts baseline because the facts about misinformation could induce participants to be more cautious in the same way that the treatment courses actively teach participants to be.

	No-course baseline	Facts baseline	Reasoning course	Emotions course	Combo course	Overall
Number of Characters (Q1)	53.73 (1.33)	54.55 (1.48)	51.98 (1.32)	55.36 (1.29)	53.74 (1.45)	53.89 (0.62)
Number of Words (Q1)	8.39 (0.22)	8.61 (0.25)	8.09 (0.22)	8.71 (0.22)	8.44 (0.24)	8.45 (0.1)
Number of Characters (Q2)	48.94 (1.45)	51.75 (1.48)	49.67 (1.61)	53.57 (1.48)	48.52 (1.37)	50.52 (0.66)
Number of Words (Q2)	8.99 (0.26)	9.49 (0.26)	9.11 (0.28)	9.79 (0.26)	8.84 (0.24)	9.25 (0.12)
Number of Characters (Q3)	45.42 (1.26)	45.62 (1.27)	44.33 (1.22)	45.43 (1.23)	43.27 (1.2)	44.81 (0.55)
Number of Words (Q3)	8.55 (0.23)	8.56 (0.23)	8.33 (0.22)	8.63 (0.23)	8.19 (0.22)	8.45 (0.1)
Number of Characters (Q4)	66.14 (1.82)	63.73 (1.29)	66.50 (1.24)	67.03 (1.29)	64.75 (1.34)	65.64 (0.63)
Number of Words (Q4)	12.24 (0.3)	11.91 (0.24)	12.41 (0.23)	12.53 (0.24)	12.14 (0.24)	12.25 (0.11)
Number of Characters (Q5)	66.17 (1.27)	67.52 (1.31)	69.12 (1.3)	70.43 (1.39)	67.10 (1.31)	68.10 (0.59)
Number of Words (Q5)	11.25 (0.23)	11.57 (0.23)	11.70 (0.23)	11.94 (0.24)	11.41 (0.23)	11.58 (0.1)

Table 6: Descriptive Statistics for Reflective Questions

Notes: Sample includes 5,313 participants who completed the follow-up survey. Standard errors are in parentheses below means. Q1-Q5 correspond to questions 1 to 5 in subsection 2.2.

the treatment courses? Third, are there interaction effects between the accuracy nudge and the treatment courses?

For the analyses in this section, it is helpful to think of our Main Experiment as two separate experiments, one in the pre-survey and one in the post-survey. In the pre-survey experiment, we can compare participants in the Accuracy Inter group who see the accuracy question before each sharing question to participants in the Accuracy After group who have not see any accuracy questions when they answer the sharing questions (see Figure 2). From this experiment, we can analyze the treatment effect of asking participants about the accuracy of a post before asking them for a sharing decision compared to not asking participants about accuracy at all.

In the second experiment in the post-survey, all participants have been exposed to accuracy questions in the pre-survey. In addition, some have seen text message courses. By analyzing outcomes at this point in the experiment, we can learn the treatment effect of the treatment courses combined with different levels of exposure to accuracy questions. Participants in the Accuracy After group again do not see the accuracy questions until after they answer the sharing questions, but because they are exposed to accuracy questions in the pre-survey, we learn from this group the treatment effect of the treatment courses combined with having seen ten accuracy questions five days earlier. Participants in the Accuracy Inter group have prior exposure to interweaved accuracy questions in the pre-survey, and again see the accuracy questions interweaved during the post-survey.

We focus on an alternative outcome in this section to evaluate the accuracy nudge. The outcome in Equation 1 is defined based on a change between pre- and post-survey behavior, but participants

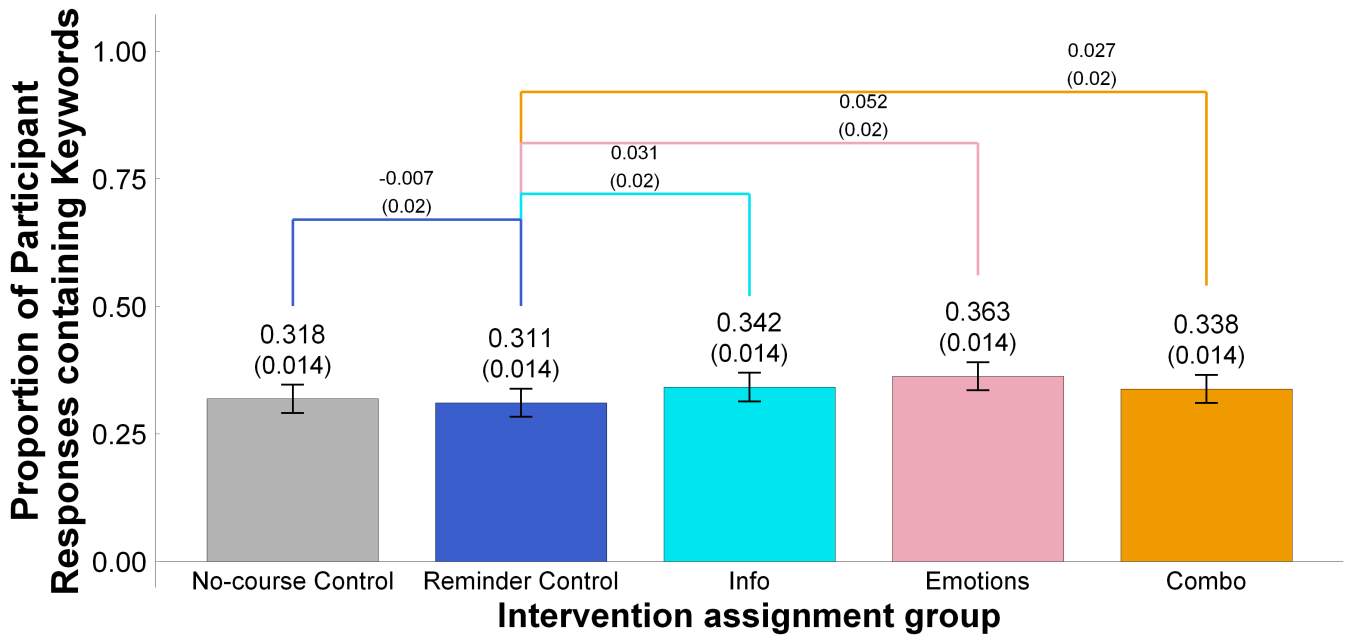


Figure 13: Self-reported Social Media Behavior (Q4) in Follow-up, by Intervention Assignment Group

Notes: Sample includes 5,313 participants who completed the follow-up survey. Each bar displays the proportion of participants whose response to question 4 in the follow-up survey contained one of the keywords. Above each bar, the standard error of each proportion is shown in parentheses below the proportion. The thin black bars represent 95% confidence intervals. Differences in the proportions are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses below.

assigned to see the accuracy nudge see the Accuracy Inter treatment in both surveys. Thus, a change in behavior between the pre- and post-survey is not a relevant metric for evaluating the nudge. We instead consider as our primary outcome the proportion of misinformation posts (out of six) a participant intends to share in the pre-survey $misinfo_i^{pre}$ or in the post-survey $misinfo_i^{post}$.

Table 5 shows that the accuracy nudge effectively decreases misinformation sharing. In the pre-survey, averaging over all text courses, participants asked to rate the accuracy of a post prior to making a sharing decision share 51.2% ($SE = 0.50$), or three of the six misinformation posts. Participants in the Accuracy After group who do not see the accuracy questions before making their sharing decision share 6.7 p.p. ($SE = 0.71$ p.p.) more misinformation, or about half a post. So, the high-dosage version of the accuracy nudge we implement changes misinformation sharing behavior in our context.

We can also evaluate the accuracy nudge using the post-survey. At this point in the experiment, participants have potentially been exposed to text message courses, and also have had previous exposure to accuracy questions (either interweaved in Accuracy Inter or at the end in Accuracy After). For simplicity, consider the No-Course baseline group. Then, the difference in post-survey sharing outcomes between Accuracy Inter and Accuracy After groups compares a group that saw interweaved accuracy five days earlier in the pre-survey, and then again in the post-survey, to a group that saw ten accuracy questions five days earlier in the pre-survey, but did not see accuracy




Figure 14: Self-reported Tip to Friend (Q5) in Follow-up, by Intervention Assignment Group

Notes: The sample is composed of the 5,313 participants who completed the follow-up survey. Each bar displays the average proportion of participants whose response to question 5 in the follow-up survey contained one of the keywords. Above each bar, the standard error of each proportion is shown in parentheses below the proportion. The thin black bars represent 95% confidence intervals. Differences in the proportions are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses below.

questions in the post-survey until they had answered all sharing questions. Table 5 shows that the Accuracy Inter group in the No-course baseline shared 51.4% ($SE = 1.22$) of misinformation posts in the post-survey relative to 57.8% ($SE = 1.11$) in the Accuracy After group, yielding a treatment effect of 6.4 p.p. ($SE = 1.65$) for the accuracy nudge, which is similar but slightly smaller than the pre-survey treatment effect. Appendix B.6 shows that qualitatively similar results hold for pre-survey and post-survey sharing and accuracy discernment.

We next compare the estimated treatment effect of the accuracy nudge to the estimated treatment effect of the treatment courses to learn whether the treatment courses are more or less effective than the accuracy nudge. Because the 18 p.p. treatment effect of the courses we estimate in subsection 3.2 uses a different outcome (Equation 1) than we use to estimate the treatment effect of the accuracy nudge in this subsection, we need to estimate an alternative treatment effect for the treatment courses using the same outcome measure: post-survey misinformation sharing.

To estimate the alternative treatment effect of the treatment courses, we restrict our sample to the 3,454 participants in the Accuracy After group and in one of the treatment courses (Emotions, Reasoning, or Combo) or the No-course baseline. Participants in this group are all exposed to ten accuracy questions in the pre-survey. Taking the difference between the post-survey misinformation sharing of participants exposed to one of the treatment courses and participants exposed to no course yields the treatment effect of the courses (in the absence of Accuracy Inter, but for participants who have previously been exposed to ten Accuracy questions in the pre-survey).

Table 5 shows that participants in the No-course baseline, Accuracy After assignment group share 57.8% ($SE = 1.11$) of the six misinformation posts in the post-survey. Participants in the Accuracy After group exposed to one of the treatment courses share 18.0 p.p. ($SE = 1.28$ p.p.) less misinformation. Note that this 18.0 p.p. treatment effect for the treatment courses (31.1% of baseline outcome) is very similar to the 18.1 p.p. (28% of baseline outcome) estimated using the Sharing Rate, our primary outcome for treatment course evaluation (baselines are also similar).

Since this treatment effect of the treatment courses is estimated for participants who were exposed to ten accuracy questions in the pre-survey (five days earlier), it is most meaningful to compare it to the treatment effect of the accuracy nudge for participants with similar prior experience. Thus, we compare the treatment course treatment effect to the treatment effect of the accuracy nudge on post-survey outcomes measured for those individuals in the No-course baseline. As described above, this effect is 6.4 p.p., much smaller than the 18.0 p.p. treatment effect for the text message courses, recalling that the control group for both treatment effects is the same at 57.8%.

Last, we test whether there is an interaction effect between the accuracy nudge and the treatment courses; that is, whether the two treatments are substitutes, complements, or have independent effects. To estimate this incremental effect, we take the difference of two differences, the first of which we already estimate above. The first difference estimates the difference between the two accuracy nudge groups when they are not exposed to a treatment course: $misinfo_i^{post}$ in the Accuracy Inter, No-course baseline group minus $misinfo_i^{post}$ in the Accuracy After, No-course baseline group. The second difference estimates the difference between the two accuracy nudge groups when they are exposed to one of the treatment courses: $misinfo_i^{post}$ in the Accuracy Inter, treatment course group minus $misinfo_i^{post}$ in the Accuracy After, treatment course group. The difference in these two differences captures whether asking accuracy interweaved in the sharing questions after educating participants about techniques to combat misinformation decreases misinformation more or less than the accuracy nudge alone.

We find a smaller difference between the accuracy nudge groups exposed to the treatment courses compared to the difference between the accuracy nudge groups not exposed to the treatment courses, suggesting that the accuracy nudge is not more valuable in combination with the treatment courses than alone. As estimated above, the difference in $misinfo_i^{post}$ between the two accuracy nudge groups in the No-course baseline is 6.4 p.p. ($SE = 1.65$ p.p.). The same difference in $misinfo_i^{post}$ between the two accuracy nudge groups exposed to the treatment courses is 4.5 p.p. ($SE = 0.90$ p.p.). The difference in differences is 1.9 p.p. ($SE = 1.88$ p.p.), so we can bound the incremental effect of the nudge and the courses at a 1.8 p.p. increase in the sharing rate, with 95% confidence.

4 Discussion

This paper evaluates three versions of a technique-based text message course on misinformation in a field experiment in Kenya to understand whether treating reasoning- or emotion-based techniques is most effective and whether the two approaches combined are complementary. We find that once participants are exposed to techniques used to manipulate emotions on social media posts and a “Stop and Question” strategy to counteract those techniques, adding content about reasoning-based

techniques has no meaningful advantage and can backfire on posts with emotion-based techniques only. Backfiring on posts with emotion-based techniques only is particularly concerning because emotional language is correlated with social media dispersion (Brady et al., 2017; Pröllochs et al., 2021). Thus, the Emotions course is more effective on the posts that spread fastest on social media. Further underscoring the importance of treating emotions, our findings are consistent with the literature drawing attention to the connection between an emotional state and belief in misinformation (Martel et al., 2020), suggesting that an intervention targeted at emotions that decreases misinformation sharing could also decrease belief in misinformation. In summary, our study demonstrates that the focus of the literature on counteracting reasoning-based techniques may be better directed in the future to learning how best to counteract emotion-based techniques.

Overall, we find that the expert-developed text message courses effectively decrease misinformation sharing both immediately after completing the course and seven to eleven weeks later. We also show that a high-dosage accuracy nudge decreases misinformation sharing, but not as much as the treatment courses. We find no evidence of complementarities from adding the accuracy nudge to the treatment courses, further confirming the preeminence of teaching about emotion-based techniques. Last, we show that participants in the treatment courses are more likely to report using (and telling friends about) strategies taught in the course than those not exposed to the treatment courses. The long-run effects of the course, its performance relative to the accuracy nudge, and the self-reported behavior in the follow-up survey all support the generalizability of our conclusions. At the same time, future research that tests these interventions using on-platform behavior, as opposed to the survey behavior we and the vast majority of the literature use, is an important next step.

Another direction for future work is to understand the precise mechanisms driving emotional sharing on social media. We propose that emotion-based techniques operates through the first-stage of a two-stage decision process in which social media users must first stop to evaluate whether a post is misinformation before they can conduct the evaluation, but several mechanisms could be at work here. For example, teaching users about emotion-based techniques may lead them to stop to evaluate a post that contains emotional language either because 1) the user intellectually identifies the emotional language and uses that as a signal for a stopping mechanism, and/or 2) the user feels strong emotions and uses that as a signal for a stopping mechanism. Our study provides some evidence for the second explanation as the Emotions course was as effective as the Reasoning or Combo course for Reasoning posts, but a research design to identify mechanisms would be better suited to answer this question.

References

- Arechar, A. A., J. N. L. Allen, a. berinsky, R. Cole, Z. Epstein, K. Garimella, A. Gully, J. G. Lu, R. M. Ross, M. Stagnaro, and et al. (2022, Feb). Understanding and reducing online misinformation across 16 countries on six continents.
- Armanasco, A. A., Y. D. Miller, B. S. Fjeldsoe, and A. L. Marshall (2017). Preventive health behavior change text message interventions: a meta-analysis. *American journal of preventive medicine* 52(3), 391–402.
- Athey, S., M. Cersosimo, K. Koutout, and Z. Li (2022, July). Inoculation against misinformation.

- Basol, M., J. Roozenbeek, M. Berriche, F. Uenal, W. P. McClanahan, and S. v. d. Linden (2021). Towards psychological herd immunity: Cross-cultural evidence for two prebunking interventions against covid-19 misinformation. *Big Data & Society* 8(1), 205395172111013868.
- Basol, M., J. Roozenbeek, and S. Van der Linden (2020). Good news about bad news: Gamified inoculation boosts confidence and cognitive immunity against fake news. *Journal of cognition* 3(1).
- Brady, W. J., J. A. Wills, J. T. Jost, J. A. Tucker, and J. J. Van Bavel (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences* 114(28), 7313–7318.
- Bryan, C. J., D. S. Yeager, C. P. Hinojosa, A. Chabot, H. Bergen, M. Kawamura, and F. Steubing (2016). Harnessing adolescent values to motivate healthier eating. *Proceedings of the National Academy of Sciences* 113(39), 10830–10835.
- Castleman, B. L. and L. C. Page (2015). Summer nudging: Can personalized text messages and peer mentor outreach increase college going among low-income high school graduates? *Journal of Economic Behavior & Organization* 115, 144–160.
- Compton, J., S. van der Linden, J. Cook, and M. Basol (2021). Inoculation theory in the post-truth era: Extant findings and new frontiers for contested science, misinformation, and conspiracy theories. *Social and Personality Psychology Compass* 15(6), e12602.
- Cook, J. (2020). Deconstructing climate science denial. *Research handbook on communicating climate change*, 62–78.
- Cook, J., S. Lewandowsky, and U. K. Ecker (2017). Neutralizing misinformation through inoculation: Exposing misleading argumentation techniques reduces their influence. *PloS one* 12(5), e0175799.
- Guess, A. M., M. Lerner, B. Lyons, J. M. Montgomery, B. Nyhan, J. Reifler, and N. Sircar (2020). A digital media literacy intervention increases discernment between mainstream and false news in the united states and india. *Proceedings of the National Academy of Sciences* 117(27), 15536–15545.
- Lewandowsky, S. and S. Van Der Linden (2021). Countering misinformation and fake news through inoculation and prebunking. *European Review of Social Psychology* 32(2), 348–384.
- List, J. A. (2020). Non est disputandum de generalizability? a glimpse into the external validity trial. Technical report, National Bureau of Economic Research.
- Maertens, R., J. Roozenbeek, M. Basol, and S. van der Linden (2021). Long-term effectiveness of inoculation against misinformation: Three longitudinal experiments. *Journal of Experimental Psychology: Applied* 27(1), 1.
- Martel, C., G. Pennycook, and D. G. Rand (2020). Reliance on emotion promotes belief in fake news. *Cognitive research: principles and implications* 5(1), 1–20.
- McGuire, W. J. (1961). Resistance to persuasion conferred by active and passive prior refutation of the same and alternative counterarguments. *The Journal of Abnormal and Social Psychology* 63(2), 326.

- McRae, K. and J. J. Gross (2020). Emotion regulation. *Emotion* 20(1), 1.
- Nguyen, T. and S. Cecchini (2021). Countering covid-19 misinformation in africa. <https://www.thinkglobalhealth.org/article/countering-covid-19-misinformation-africa>.
- Offer-Westort, M., L. R. Rosenzweig, and S. Athey (2021). Optimal policies to battle the coronavirus “infodemic” among social media users in sub-saharan africa. *OSF Registered Study*.
- Pennycook, G., Z. Epstein, M. Mosleh, A. A. Arechar, D. Eckles, and D. G. Rand (2021). Shifting attention to accuracy can reduce misinformation online. *Nature* 592(7855), 590–595.
- Petty, R. and L. Hinsenkamp (2017). The sage encyclopedia of political behavior.
- Petty, R. E. and P. Briñol (2015). Emotion and persuasion: Cognitive and meta-cognitive processes impact attitudes. *Cognition and Emotion* 29(1), 1–26.
- Petty, R. E. and J. T. Cacioppo (1986). The elaboration likelihood model of persuasion. In *Communication and persuasion*, pp. 1–24. Springer.
- Pröllochs, N., D. Bär, and S. Feuerriegel (2021). Emotions explain differences in the diffusion of true vs. false social media rumors. *Scientific reports* 11(1), 1–12.
- Rodgers, A., T. Corbett, D. Bramley, T. Riddell, M. Wills, R.-B. Lin, and M. Jones (2005). Do u smoke after txt? results of a randomised trial of smoking cessation using mobile phone text messaging. *Tobacco control* 14(4), 255–261.
- Romano, J. P. and M. Wolf (2007). Control of generalized error rates in multiple testing. *The Annals of Statistics* 35(4), 1378–1408.
- Roozenbeek, J., R. Maertens, W. McClanahan, and S. van der Linden (2021). Disentangling item and testing effects in inoculation research on online misinformation: Solomon revisited. *Educational and Psychological Measurement* 81(2), 340–362.
- Roozenbeek, J. and S. Van Der Linden (2019). The fake news game: actively inoculating against the risk of misinformation. *Journal of risk research* 22(5), 570–580.
- Roozenbeek, J., S. Van Der Linden, B. Goldberg, S. Rathje, and S. Lewandowsky (2022). Psychological inoculation improves resilience against misinformation on social media. *Science advances* 8(34), eabo6254.
- Suffoletto, B., J. Kristan, C. Callaway, K. H. Kim, T. Chung, P. M. Monti, and D. B. Clark (2014). A text message alcohol intervention for young adult emergency department patients: a randomized clinical trial. *Annals of emergency medicine* 64(6), 664–672.
- Telzer, E. H., A. J. Fuligni, M. D. Lieberman, and A. Galván (2014). Neural sensitivity to eudaimonic and hedonic rewards differentially predict adolescent depressive symptoms over time. *Proceedings of the National Academy of Sciences* 111(18), 6600–6605.
- van der Linden, S. (2022). Misinformation: susceptibility, spread, and interventions to immunize the public. *Nature Medicine* 28(3), 460–467.
- Van der Linden, S., A. Leiserowitz, S. Rosenthal, and E. Maibach (2017). Inoculating the public against misinformation about climate change. *Global Challenges* 1(2), 1600008.

York, B. N., S. Loeb, and C. Doss (2019). One step at a time the effects of an early literacy text-messaging program for parents of preschoolers. *Journal of Human Resources* 54(3), 537–566.

A Tests with Multiple Hypothesis Test Correction

In this appendix, we report the execution of tests from our pre-analysis plan and tests on our non-prespecified Follow-up Experiment. Since we execute multiple tests to evaluate the same hypothesis (e.g., “Do the treatment courses decrease misinformation sharing?” and “Which treatment course is most effective at decreasing misinformation sharing?”), we want to control the probability of drawing at least one false conclusion.

To control the Family-Wise Error Rate (FWER), we correct for multiple hypothesis testing by using the Romano-Wolf procedure (Romano and Wolf, 2007) to adjust p-values. These adjusted p-values are an extension of resample-based p-values for single hypothesis testing, which are defined as the fraction of resamples that produce a Studentized null statistic with a value that is more extreme than the test statistic from the test with the original sample. The Romano-Wolf p-values are defined as the fraction of resamples that produce a “max”-statistic that is more extreme than the test statistic from the test with the original sample, where the “max”-statistic is defined, for a given test T , as the maximum of the Studentized null statistics from the set of Studentized null statistics corresponding to T and the tests for which the statistic with the original data are less extreme than the statistic from T . By capturing the correlation between test statistics through the use of the “max”-statistics, this procedure achieves greater power than other corrections such as Bonferroni or Holm.

We group our hypotheses into five families for this correction. The first family of hypotheses tests whether the treatment courses (Emotions, Reasoning, and Combo) are more effective than the baselines (No-course and Facts). The second family of hypotheses tests which of the treatment courses works best and which work best for each type of misinformation post. The third family of hypotheses tests the effects of the accuracy nudge, separately, and then relative to and in combination with the treatment courses. The fourth family of hypotheses tests whether the treatment courses are more effective than the Facts baseline in the follow-up survey to test the long-run effects of the treatment courses. Last, the fifth family of hypotheses tests which of the treatment courses work best in the long-run, and which work best for each type of misinformation post in the follow-up survey.

A.1 Family 1 Tests

Group 1	Group 2	N. Obs. 1	N. Obs. 2	Mean 1	Mean 2	Diff. in Means	95% CI	Unadj p-Value	RW Adj p-Value
All	No-course	4,646	1,454	0.4594 (0.0059)	0.6390 (0.0101)	-0.1797 (0.0117)	$(-\infty, -0.1605)$	0.0000	0.0010
All	Facts	4,646	1,584	0.4594 (0.0059)	0.5633 (0.0103)	-0.1040 (0.0118)	$(-\infty, -0.0845)$	0.0000	0.0010
Reasoning	No-course	1,502	1,454	0.4963 (0.0103)	0.6390 (0.0101)	-0.1427 (0.0144)	$(-\infty, -0.119)$	0.0000	0.0010
Emotions	No-course	1,617	1,454	0.4301 (0.0098)	0.6390 (0.0101)	-0.2089 (0.0141)	$(-\infty, -0.1858)$	0.0000	0.0010
Combo	No-course	1,527	1,454	0.4540 (0.0105)	0.6390 (0.0101)	-0.1850 (0.0145)	$(-\infty, -0.1611)$	0.0000	0.0010
Reasoning	Facts	1,502	1,584	0.4963 (0.0103)	0.5633 (0.0103)	-0.0670 (0.0146)	$(-\infty, -0.043)$	0.0000	0.0010
Emotions	Facts	1,617	1,584	0.4301 (0.0098)	0.5633 (0.0103)	-0.1332 (0.0142)	$(-\infty, -0.1098)$	0.0000	0.0010
Combo	Facts	1,527	1,584	0.4540 (0.0105)	0.5633 (0.0103)	-0.1093 (0.0147)	$(-\infty, -0.0852)$	0.0000	0.0010
Facts	No-course	1,584	1,454	0.5633 (0.0103)	0.6390 (0.0101)	-0.0757 (0.0144)	$(-\infty, -0.0521)$	0.0000	0.0010

Table A1: Average Treatment Effects

Notes: The sample is composed of the 7,774 participants who completed the post-survey and shared at least one non-misinformation post in the pre-survey, pooling participants in the Accuracy Inter and Accuracy After treatments. Each row displays the results of one-sided t-tests. For the assignment groups specified in the first column of the table, the t-tests compare the ratio of the number of misinformation posts the participant shared in the post-survey that correspond to the non-misinformation counterpart (non-misinformation post) they shared in the pre-survey to the number of those corresponding non-misinformation posts they shared in the pre-survey. Romano-Wolf adjusted p-values are computed jointly for all Family 1 tests using 1000 simulations and therefore have a minimum of 0.001.

Post Shared	Group 1	Group 2	N. Obs. 1	N. Obs. 2	Mean 1	Mean 2	Diff. in Means	95% CI	Unadj p-Value	RW Adj p-Value
1	All	No-course	786	237	0.3359 (0.0169)	0.4515 (0.0324)	-0.1156 (0.0365)	($-\infty$, -0.0554)	0.0008	0.0110
2	All	No-course	1,427	453	0.3980 (0.0101)	0.5861 (0.0177)	-0.1881 (0.0204)	($-\infty$, -0.1545)	0.0000	0.0010
3	All	No-course	2,433	764	0.5353 (0.0076)	0.7286 (0.0113)	-0.1933 (0.0136)	($-\infty$, -0.1709)	0.0000	0.0010
1	All	Facts	786	283	0.3359 (0.0169)	0.4064 (0.0292)	-0.0705 (0.0338)	($-\infty$, -0.0148)	0.0187	0.0619
2	All	Facts	1,427	519	0.3980 (0.0101)	0.4875 (0.0179)	-0.0894 (0.0205)	($-\infty$, -0.0557)	0.0000	0.0010
3	All	Facts	2,433	782	0.5353 (0.0076)	0.6705 (0.0123)	-0.1352 (0.0144)	($-\infty$, -0.1115)	0.0000	0.0010
1	Reasoning	No-course	236	237	0.4025 (0.0320)	0.4515 (0.0324)	-0.0489 (0.0455)	($-\infty$, 0.0261)	0.1415	0.3177
2	Reasoning	No-course	466	453	0.4303 (0.0177)	0.5861 (0.0177)	-0.1558 (0.0250)	($-\infty$, -0.1146)	0.0000	0.0010
3	Reasoning	No-course	800	764	0.5625 (0.0130)	0.7286 (0.0113)	-0.1661 (0.0172)	($-\infty$, -0.1377)	0.0000	0.0010
1	Emotions	No-course	280	237	0.3000 (0.0274)	0.4515 (0.0324)	-0.1515 (0.0425)	($-\infty$, -0.0815)	0.0002	0.0080
2	Emotions	No-course	497	453	0.3853 (0.0166)	0.5861 (0.0177)	-0.2008 (0.0243)	($-\infty$, -0.1608)	0.0000	0.0010
3	Emotions	No-course	840	764	0.5000 (0.0128)	0.7286 (0.0113)	-0.2286 (0.0171)	($-\infty$, -0.2005)	0.0000	0.0010
1	Combo	No-course	270	237	0.3148 (0.0283)	0.4515 (0.0324)	-0.1367 (0.0430)	($-\infty$, -0.0658)	0.0008	0.0100
2	Combo	No-course	464	453	0.3793 (0.0180)	0.5861 (0.0177)	-0.2068 (0.0252)	($-\infty$, -0.1653)	0.0000	0.0010
3	Combo	No-course	793	764	0.5452 (0.0134)	0.7286 (0.0113)	-0.1834 (0.0176)	($-\infty$, -0.1545)	0.0000	0.0010
1	Reasoning	Facts	236	283	0.4025 (0.0320)	0.4064 (0.0292)	-0.0038 (0.0433)	($-\infty$, 0.0676)	0.4649	0.4715
2	Reasoning	Facts	466	519	0.4303 (0.0177)	0.4875 (0.0179)	-0.0572 (0.0251)	($-\infty$, -0.0158)	0.0116	0.0589
3	Reasoning	Facts	800	782	0.5625 (0.0130)	0.6705 (0.0123)	-0.1080 (0.0179)	($-\infty$, -0.0786)	0.0000	0.0010
1	Emotions	Facts	280	283	0.3000 (0.0274)	0.4064 (0.0292)	-0.1064 (0.0401)	($-\infty$, -0.0403)	0.0041	0.0270
2	Emotions	Facts	497	519	0.3853 (0.0166)	0.4875 (0.0179)	-0.1022 (0.0244)	($-\infty$, -0.062)	0.0000	0.0010
3	Emotions	Facts	840	782	0.5000 (0.0128)	0.6705 (0.0123)	-0.1705 (0.0177)	($-\infty$, -0.1413)	0.0000	0.0010
1	Combo	Facts	270	283	0.3148 (0.0283)	0.4064 (0.0292)	-0.0915 (0.0407)	($-\infty$, -0.0245)	0.0125	0.0589
2	Combo	Facts	464	519	0.3793 (0.0180)	0.4875 (0.0179)	-0.1082 (0.0253)	($-\infty$, -0.0665)	0.0000	0.0010
3	Combo	Facts	793	782	0.5452 (0.0134)	0.6705 (0.0123)	-0.1253 (0.0182)	($-\infty$, -0.0954)	0.0000	0.0010
1	Facts	No-course	283	237	0.4064 (0.0292)	0.4515 (0.0324)	-0.0451 (0.0436)	($-\infty$, 0.0268)	0.1509	0.3177
2	Facts	No-course	519	453	0.4875 (0.0179)	0.5861 (0.0177)	-0.0986 (0.0251)	($-\infty$, -0.0572)	0.0000	0.0030
3	Facts	No-course	782	764	0.6705 (0.0123)	0.7286 (0.0113)	-0.0581 (0.0167)	($-\infty$, -0.0306)	0.0003	0.0080

Table A2: Treatment effects by number of non-misinformation posts shared in the pre-survey

Notes: The sample is composed of the 7,774 participants who completed the post-survey and shared at least one non-misinformation post in the pre-survey, pooling participants in the Accuracy After and Accuracy Inter treatments. The participants are divided into sharing groups according to the number of non-misinformation posts shared in the pre-survey. Number of participants per group: 1 post shared: 1306; 2 posts shared: 2399; 3 posts

A.2 Family 2 Tests

Group 1	Group 2	N. Obs. 1	N. Obs. 2	Mean 1	Mean 2	Diff. in Means	95% CI	Unadj p-Value	RW Adj p-Value
All Posts									
Reasoning	Emotions	1,502	1,617	0.4963 (0.0103)	0.4301 (0.0098)	0.0662 (0.0143)	(0.0383, 0.0942)	0.0000	0.0010
Reasoning	Combo	1,502	1,527	0.4963 (0.0103)	0.4540 (0.0105)	0.0423 (0.0147)	(0.0135, 0.0711)	0.0041	0.0320
Emotions	Combo	1,617	1,527	0.4301 (0.0098)	0.4540 (0.0105)	-0.0239 (0.0144)	(-0.0521, 0.0042)	0.0956	0.3297
Reasoning Posts									
Reasoning	Emotions	1,181	1,261	0.5394 (0.0145)	0.5091 (0.0141)	0.0303 (0.0202)	(-0.0094, 0.0699)	0.1347	0.3716
Reasoning	Combo	1,181	1,200	0.5394 (0.0145)	0.4900 (0.0144)	0.0494 (0.0205)	(0.0092, 0.0895)	0.0159	0.0789
Emotions	Combo	1,261	1,200	0.5091 (0.0141)	0.4900 (0.0144)	0.0191 (0.0202)	(-0.0204, 0.0587)	0.3432	0.4785
Emotions Posts									
Emotions	Reasoning	1,254	1,192	0.4537 (0.0141)	0.5327 (0.0145)	-0.0790 (0.0202)	($-\infty$, -0.0458)	0.0000	0.0010
Emotions	Combo	1,254	1,193	0.4537 (0.0141)	0.5105 (0.0145)	-0.0567 (0.0202)	($-\infty$, -0.0235)	0.0025	0.0330
Reasoning	Combo	1,192	1,193	0.5327 (0.0145)	0.5105 (0.0145)	0.0222 (0.0205)	(-0.0179, 0.0624)	0.2772	0.4785
Combo Posts									
Combo	Reasoning	1,184	1,195	0.4535 (0.0145)	0.4803 (0.0145)	-0.0268 (0.0205)	($-\infty$, 0.0069)	0.0953	0.4056
Combo	Emotions	1,184	1,279	0.4535 (0.0145)	0.4034 (0.0137)	0.0501 (0.0199)	(0.011, 0.0892)	0.0121	0.0649
Reasoning	Emotions	1,195	1,279	0.4803 (0.0145)	0.4034 (0.0137)	0.0769 (0.0199)	(0.0378, 0.116)	0.0001	0.0010

Table A3: Differences in sharing by type of post

Notes: The sample is composed of the 4,646 participants in the Reasoning, Emotions or Combo intervention assignment groups who completed the post-survey and shared at least one non-misinformation post in the pre-survey, pooling participants in the Accuracy After and Accuracy Inter treatments. Each row displays the results of t-tests. For the assignment groups specified in the first column of the table, the t-tests compare the ratio of the number of misinformation posts the participant shared in the post-survey that correspond to the non-misinformation counterpart (non-misinformation post) they shared in the pre-survey to the number of those corresponding non-misinformation posts they shared in the pre-survey. Tests 10, 11, 12, 13, 14, 18, 20 and 21 are two-sided t-tests. Tests 16, 17 and 19 are one-sided t-tests. Romano-Wolf adjusted p-values are computed jointly for all Family 2 tests using 1000 simulations and therefore have a minimum of 0.001.

A.3 Family 3 Tests

Group 1	Group 2	N. Obs. 1	N. Obs. 2	Mean 1	Mean 2	Diff. in Means	95% CI	Unadj p-Value	RW Adj p-Value
Pre									
Accuracy Inter	Accuracy After	4,357	4,321	0.5124 (0.0052)	0.5796 (0.0047)	-0.0672 (0.0071)	$(-\infty, -0.0556)$	0.0000	0.0010
No-course Baseline Post									
Accuracy Inter	Accuracy After	800	826	0.5138 (0.0122)	0.5779 (0.0111)	-0.0641 (0.0165)	$(-\infty, -0.037)$	0.0001	0.0010

Table A4: Effect of accuracy nudge

Notes: The sample is composed of the 8,768 participants who completed the post-survey. Each row displays the results of one-sided t-tests comparing the average of the proportion of the shared misinformation posts in the pre- or post-survey, respectively, for the assignment groups specified in the rows in bold. Romano-Wolf adjusted p-values are computed jointly for all Family 3 tests using 1000 simulations and therefore have a minimum of 0.001.

Group	N Obs.	Mean Pre	Mean Post	Mean Diff.	95% CI	Unadj p-Value	Holm Adj p-Value
Accuracy After	2,654	0.571 (0.006)	0.398 (0.006)	-0.1728 (0.0650)	$(-\infty, -0.1620)$	0.0000	0.0000
Accuracy Inter	2,609	0.514 (0.007)	0.353 (0.006)	-0.1604 (0.0066)	$(-\infty, -0.1495)$	0.0000	0.0000
Diff in Mean Diff	5,263	-	-	-0.0124 (0.0093)	$(-0.0277, 0.0029)$	0.1816	0.1816

Table A5: Effect of courses and accuracy nudge

Notes: The sample is composed of the 5,263 participants in the Reasoning, Emotions, and Combo assignment groups who completed the post-survey. The first two rows display the results of paired one-sided t-tests that compare the post-pre difference in the proportion of shared misinformation posts for the Accuracy After and Accuracy Inter groups, respectively. The last row displays the results of the test of the difference-in-mean-differences coefficient being equal to 0. Romano-Wolf adjusted p-values are computed jointly for all Family 3 tests using 1000 simulations and therefore have a minimum of 0.001.

Group 1	Group 2	N. Obs. 1	N. Obs. 2	Mean 1	Mean 2	Diff. in Means	95% CI	Unadj p-Value	RW Adj p-Value
Treatment After Post	All Inter Post	2,609	4,357	0.3984 (0.0064)	0.5124 (0.0052)	-0.1141 (0.0083)	(-0.1303, -0.0978)	0.0000	0.0010
Treatment After Post	No-course Inter Post	2,609	800	0.3984 (0.0064)	0.5138 (0.0122)	-0.1154 (0.0138)	(-0.1424, -0.0884)	0.0000	0.0010

Table A6: Misinformation Sharing: Treatment Courses vs Accuracy Nudge

Notes: The sample is composed of the 8,768 participants who completed the post-survey. Each row displays the results of two-sided t-tests comparing the average of the proportion of shared misinformation posts in the pre- or post-survey, respectively, for the assignment groups specified in the rows in bold. Romano-Wolf adjusted p-values are computed jointly for all Family 3 tests using 1000 simulations and therefore have a minimum of 0.001.

A.4 Family 4 Tests

Group 1	Group 2	N. Obs. 1	N. Obs. 2	Mean 1	Mean 2	Diff. in Means	95% CI	Unadj p-Value	RW Adj p-Value
All	Facts	2,554	857	0.4161 (0.0081)	0.5068 (0.0140)	-0.0907 (0.0162)	$(-\infty, -0.064)$	0.0000	0.0010
Reasoning	Facts	857	857	0.4389 (0.0139)	0.5068 (0.0140)	-0.0679 (0.0198)	$(-\infty, -0.0353)$	0.0003	0.0010
Emotions	Facts	865	857	0.3840 (0.0140)	0.5068 (0.0140)	-0.1228 (0.0198)	$(-\infty, -0.0902)$	0.0000	0.0010
Combo	Facts	832	857	0.4259 (0.0142)	0.5068 (0.0140)	-0.0809 (0.0200)	$(-\infty, -0.0481)$	0.0000	0.0010
Facts	No-course	857	838	0.5068 (0.0140)	0.4873 (0.0142)	0.0195 (0.0200)	$(-0.0197, 0.0588)$	0.3289	0.3447

Table A7: Average Treatment Effects - Follow-up

Notes: The sample is composed of the 4,249 participants who completed the follow-up survey and shared at least one non-misinformation post in the post-survey, pooling participants in the Accuracy After and Accuracy Inter treatments. Each row displays the results of t-tests. For the assignment groups specified in the first column of the table, the t-tests compare the ratio of the number of misinformation posts the participant shares in the follow-up survey that correspond to the non-misinformation counterpart (non-misinformation post) they shared in the post-survey to the number of those corresponding non-misinformation posts they shared in the post-survey. Test 9F is two-sided. Tests 2F, 6F, 7F and 8F are one-sided. Romano-Wolf adjusted p-values are computed jointly for all Family 4 tests using 1000 simulations and therefore have a minimum of 0.001.

A.5 Family 5 Tests

Group 1	Group 2	N. Obs. 1	N. Obs. 2	Mean 1	Mean 2	Diff. in Means	95% CI	Unadj p-Value	RW Adj p-Value
All Posts									
Reasoning	Emotions	857	865	0.4389 (0.0139)	0.3840 (0.0140)	0.0549 (0.0197)	(0.0162, 0.0936)	0.0054	0.0470
Reasoning	Combo	857	832	0.4389 (0.0139)	0.4259 (0.0142)	0.0131 (0.0199)	(-0.026, 0.0521)	0.5118	0.9471
Emotions	Combo	865	832	0.3840 (0.0140)	0.4259 (0.0142)	-0.0419 (0.0199)	(-0.0809, -0.0028)	0.0355	0.2218
Reasoning Posts									
Reasoning	Emotions	599	587	0.5159 (0.0204)	0.4668 (0.0206)	0.0491 (0.0290)	(-0.0079, 0.106)	0.0911	0.4266
Reasoning	Combo	599	553	0.5159 (0.0204)	0.4955 (0.0213)	0.0204 (0.0295)	(-0.0375, 0.0783)	0.4899	0.9471
Emotions	Combo	587	553	0.4668 (0.0206)	0.4955 (0.0213)	-0.0287 (0.0296)	(-0.0868, 0.0294)	0.3329	0.8761
Emotions Posts									
Emotions	Reasoning	585	594	0.4188 (0.0204)	0.4983 (0.0205)	-0.0795 (0.0290)	($-\infty$, -0.0318)	0.0031	0.0509
Emotions	Combo	585	587	0.4188 (0.0204)	0.4889 (0.0206)	-0.0701 (0.0290)	($-\infty$, -0.0223)	0.0079	0.1089
Reasoning	Combo	594	587	0.4983 (0.0205)	0.4889 (0.0206)	0.0094 (0.0291)	(-0.0477, 0.0665)	0.7472	0.9471
Combo Posts									
Combo	Reasoning	564	582	0.4167 (0.0208)	0.4158 (0.0204)	0.0009 (0.0292)	(-0.0563, 0.0581)	0.9765	0.9471
Combo	Emotions	564	579	0.4167 (0.0208)	0.3990 (0.0204)	0.0177 (0.0291)	(-0.0394, 0.0748)	0.5430	0.9471
Reasoning	Emotions	582	579	0.4158 (0.0204)	0.3990 (0.0204)	0.0168 (0.0289)	(-0.0398, 0.0735)	0.5596	0.9471

Table A8: Differences in sharing by type of post - Follow-up

Notes: The sample is composed of the 2,554 participants in the Reasoning, Emotions or Combo intervention assignment groups who completed the follow-up survey and shared at least one non-misinformation post in the post survey, pooling participants in the Accuracy Inter and Accuracy After treatments. Each row displays the results of t-tests. The t-tests compare the ratio of the number of misinformation posts the participant shares in the follow-up survey that correspond to the non-misinformation counterpart (non-misinformation post) they shared in the post-survey to the number of those corresponding non-misinformation posts they did share in the post-survey, for the assignment groups specified in the first column of the table, for the assignment groups specified in the first column of the table, and separately by type of post. Tests 10F, 11F, 12F, 13F, 14F, 15F, 18F, 19F, 20F and 21F are two-sided t-tests. Tests 16F and 17F are one-sided t-tests. Romano-Wolf adjusted p-values are computed jointly for all Family 5 tests using 1000 simulations and therefore have a minimum of 0.001.

B Additional Tables and Figures

B.1 Ad Cost and Text Message Course Funnel Statistics

	Count	% of Previous Funnel Stage	Cost Per
Impressions	19,640,962	-	\$0.001
Clicks	193,227	0.98%	\$0.054
Started Pre-survey	25287	13.08%	\$0.41
Completed Pre-survey	22526	89.08%	\$0.46
Started Text Message Course	18598	82.56%	\$0.56
Completed Day 1 Course	16684	89.71%	\$0.63
Completed Day 2 Course	13997	83.89%	\$0.75
Completed Day 3 Course	13093	93.54%	\$0.80
Completed Day 4 Course	11396	87.04%	\$0.92
Completed Entire Course	10934	95.95%	\$0.96
Started Post-survey	9589	87.70%	\$1.09
Completed Post-survey	8678	90.50%	\$1.21
Started Follow-up	5783	66.64%	\$1.81
Completed Follow-up	5313	91.87%	\$1.97

Table B1: Full Funnel with Ad Costs

Notes: The ads metrics, Impressions and Clicks, are extracted from Facebook Ads manager. The total cost of ads was \$10,460.26. The number of participants who started the pre-survey is an upper bound estimate because we could count only survey copies on Qualtrics and not identify users. We were only able to identify users once they completed the pre-survey and provided a phone number, which is how we ensured that all down-funnel outcomes counted unique users based on phone number. The post- and follow-up surveys required users to validate their phone number before they could start the survey. We received 40,845 survey copies in total for the pre-survey, of which we discarded 3,092 users who had participated in one of our pilot studies and another 12,466 survey copies filled out by duplicated phone numbers. In the post-survey, we filtered out participants who encountered system errors ($N = 104$), did not have at least 5 days in between pre- and post-survey dates ($N = 2,101$), and/or did not complete the full text-message course ($N = 509$). For the follow-up survey, we filtered out 2,371 survey copies filled out by duplicated phone numbers and retained only the first copy done by each phone number.

	No-course baseline	Facts baseline	Reasoning course	Emotions course	Combo course
Started Text Message Course	1646	3746	3686	3755	3713
	-	-	-	-	-
Completed Day 1 Course	1353 (82.2%)	3276 (87.45%)	3234 (87.74%)	3240 (86.28%)	3236 (87.15%)
Completed Day 2 Course	1122 (68.17%)	2592 (69.19%)	2575 (69.86%)	2551 (67.94%)	2581 (69.51%)
Completed Day 3 Course	1029 (62.52%)	2356 (62.89%)	2260 (61.31%)	2373 (63.2%)	2407 (64.83%)
Completed Day 4 Course	927 (56.32%)	2195 (58.6%)	2060 (55.89%)	2225 (59.25%)	2194 (59.09%)
Completed Entire Course	856 (52%)	2090 (55.79%)	1960 (53.17%)	2124 (56.56%)	2051 (55.24%)

Table B2: Text Message Course Funnel, by Assignment Group

Notes: The numbers in percent (in brackets) represent the % of participants who started the text message course in the respective assignment groups.

B.2 Covariates Explanation

Covariate in Survey	Constructed Covariate	Explanation
Age	Age	Integer denoting age of participants
Gender	Man	Complement to "woman" and "other"
Education	High school or less	including "Less than a high school diploma" and "High school degree or equivalent"
	Some college	including "Some college, no degree" and "Associate degree"
	Bachelor's degree	Bachelor's degree
	Graduate degrees	including "Master's degree" and "Doctorate or professional degree"
Marital status	Married	Complement to "Single," "Widowed," "Divorced," and "Separated"
Employment status	Employed	including "Employed full time," "Employed part time," and "Self-employed"
	Unemployed	including "Unemployed and currently looking for work," "Unemployed not currently looking for work," "Retired," "Homemaker," and "Unable to work"
Location	Student	Student
	Mostly urban	Live in mostly urban area
	Suburban	live in suburban area
Religion	Mostly rural	live in mostly rural area
	Christian	Complement to "None," "Hinduism," "Muslim," "Traditionalist," and "Other"
Religiosity	Attends religious services	Frequency of attending religious services including "Less than once a month," "One to three times per month," "Once a week," "More than once a week but less than daily," and "Daily"; Complement to "Never"
Social media user	Uses social media	1 if "Yes" and 0 if "No"
Hrs/day on social media	Hrs/day on social media	Integer denoting how many hours are spent on average each day on social media
Prop. of content shared	80 - 100%	Proportion of the content the participant sees on social media that they choose to share is between 80 and 100%
	60 - 80%	Similar as above
	40 - 60%	Similar as above
	20 - 40%	Similar as above
	0 - 20%	Similar as above

Table B3: Covariates Explanation

B.3 Attrition and Balance Figures




Figure B1: Differences in Samples, by Survey Completion

Notes: The Pre sample is composed of the 13,848 participants who completed the pre-survey but not the post-survey. The Post sample is composed of the 3,365 participants who completed both the pre-survey and the post-survey, but not the follow-up. The Followup sample is composed of the 5,313 participants who completed all of the surveys. The differences are taken by subtracting the Pre sample from the Post sample, and the Post sample from the Follow-up sample.




Figure B2: Difference in Samples of Post-survey Completers, by Intervention Assignment Group

Notes: The sample is composed of the 8,768 participants who completed the post-survey. The differences are taken by subtracting the No-course baseline from each of the other groups.




Figure B3: Difference in Samples of Follow-up Survey Completers, by Intervention Assignment Group

Notes: The sample is composed of the 5,313 participants who completed the post-survey. The differences are taken by subtracting the Facts baseline from each of the other groups.

B.4 Main Results, by Accuracy Nudge Group

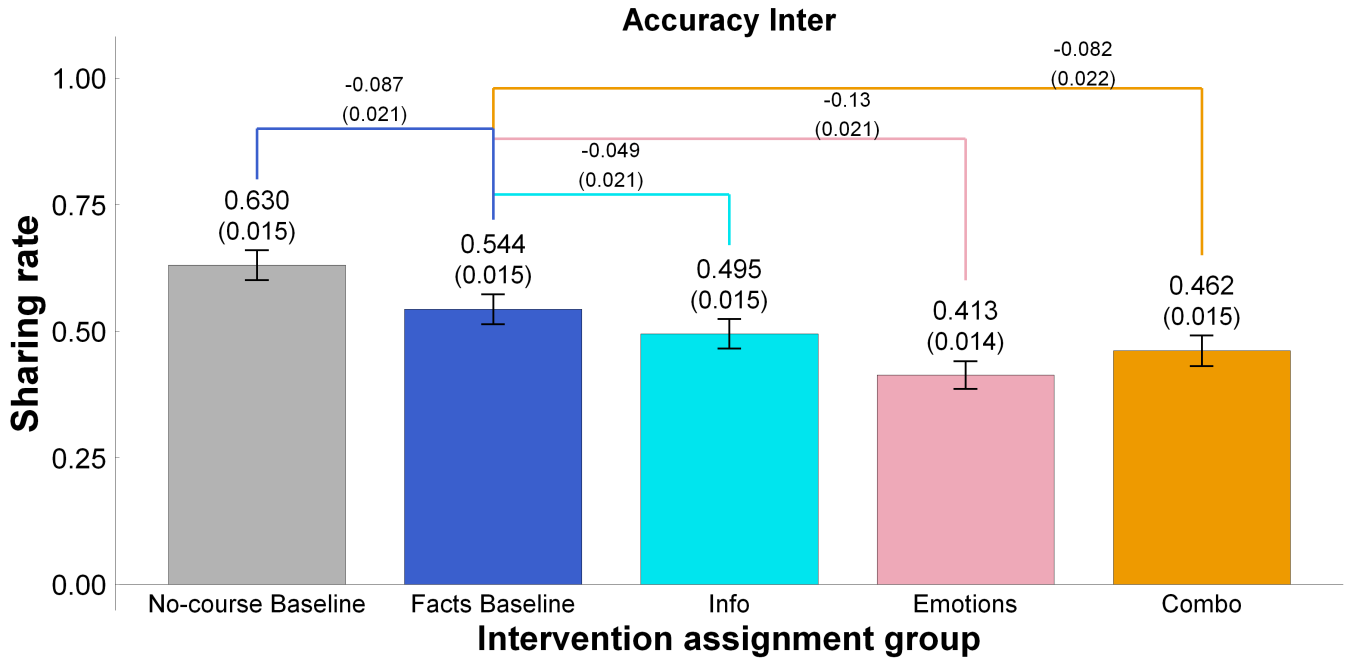


Figure B4: Misinformation Sharing, by Intervention Assignment Group

Notes: Sample includes 3,701 participants from the Accuracy Inter group who completed the post-survey and shared at least one non-misinformation post in the pre-survey. Each bar displays the Sharing Rate for misinformation posts, as defined in Equation 1, by participants in their respective intervention assignment group. Above each bar, the standard error is shown in parentheses below the Sharing Rate. The thin black bars represent 95% confidence intervals. Differences in Sharing Rates are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

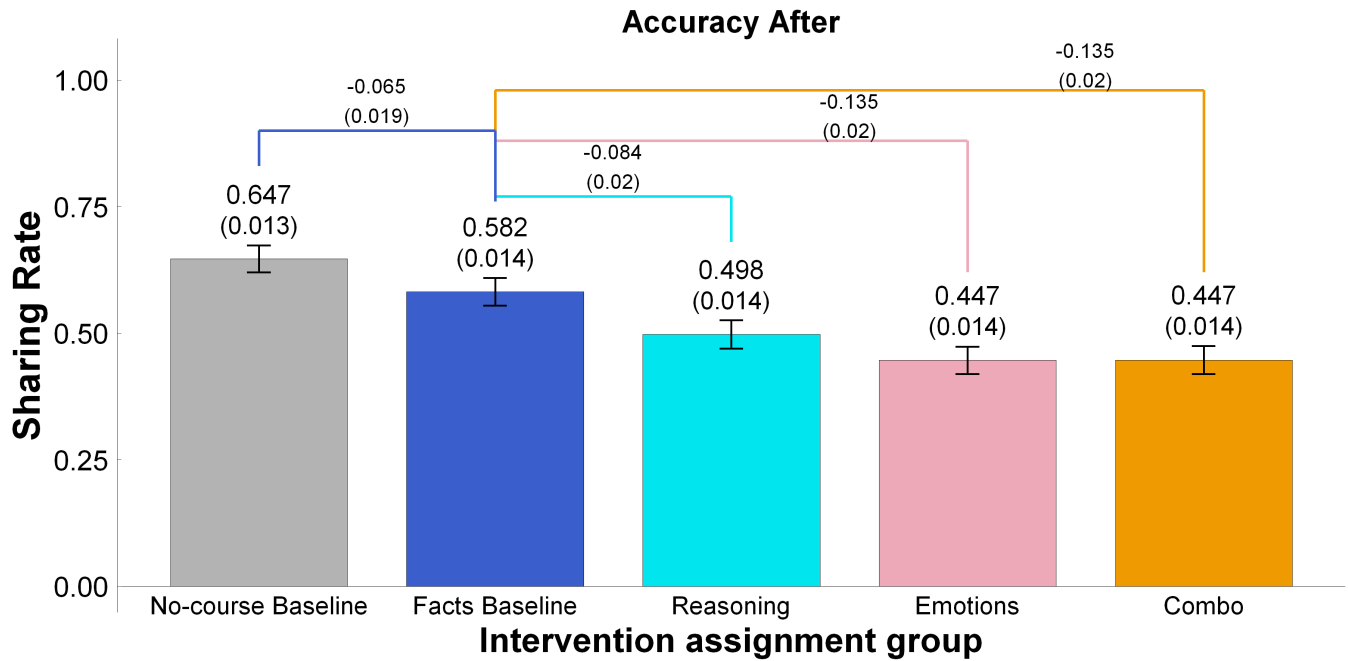


Figure B5: Misinformation Sharing, by Intervention Assignment Group

Notes: Sample includes 3,983 participants from the Accuracy After group who completed the post-survey and shared at least one non-misinformation post in the pre-survey. Each bar displays the Sharing Rate for misinformation posts, as defined in Equation 1, by participants in their respective intervention assignment group. Above each bar, the standard error is shown in parentheses below the Sharing Rate. The thin black bars represent 95% confidence intervals. Differences in Sharing Rates are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

B.5 Alternative Outcomes Family 1 Tests

Group 1	Group 2	N. Obs. 1	N. Obs. 2	Mean 1	Mean 2	Diff. in Means	95% CI	Unadj p-Value	Holm Adj p-Value
All	No-course	5,263	1,626	0.5115 (0.0265)	0.0424 (0.0465)	0.4691 (0.0536)	(0.3809, ∞)	0.0000	0.0000
All	Facts	5,263	1,789	0.5115 (0.0265)	0.2146 (0.0438)	0.2969 (0.0512)	(0.2126, ∞)	0.0000	0.0000
Info	No-course	1,689	1,626	0.3541 (0.0467)	0.0424 (0.0465)	0.3116 (0.0659)	(0.2031, ∞)	0.0000	0.0000
Emotions	No-course	1,836	1,626	0.6171 (0.0447)	0.0424 (0.0465)	0.5747 (0.0645)	(0.4685, ∞)	0.0000	0.0000
Combo	No-course	1,738	1,626	0.5529 (0.0463)	0.0424 (0.0465)	0.5105 (0.0656)	(0.4026, ∞)	0.0000	0.0000
Info	Facts	1,689	1,789	0.3541 (0.0467)	0.2146 (0.0438)	0.1394 (0.0640)	(0.034, ∞)	0.0148	0.0296
Emotions	Facts	1,836	1,789	0.6171 (0.0447)	0.2146 (0.0438)	0.4025 (0.0626)	(0.2994, ∞)	0.0000	0.0000
Combo	Facts	1,738	1,789	0.5529 (0.0463)	0.2146 (0.0438)	0.3383 (0.0637)	(0.2335, ∞)	0.0000	0.0000
Facts	No-course	1,789	1,626	0.2146 (0.0438)	0.0424 (0.0465)	0.1722 (0.0639)	(0.0671, ∞)	0.0035	0.0106
Info	Emotions	1,689	1,836	0.3541 (0.0467)	0.6171 (0.0447)	-0.2630 (0.0647)	(-0.3899, -0.1362)	0.0000	0.0002
Info	Combo	1,689	1,738	0.3541 (0.0467)	0.5529 (0.0463)	-0.1989 (0.0657)	(-0.3278, -0.07)	0.0025	0.0100
Emotions	Combo	1,836	1,738	0.6171 (0.0447)	0.5529 (0.0463)	0.0642 (0.0643)	(-0.062, 0.1903)	0.3187	0.3187

Table B4: Average Treatment Effects - Discernment Score

Notes: The sample is composed of the 8,768 participants who completed the post-survey, pooling participants in the Accuracy Inter and Accuracy After treatments. Each row displays the results of t-tests. For the assignment groups specified in the first column of the table, the t-tests compare the average of the post-pre difference in the discernment score, defined as the number of non-misinformation posts shared minus the number of misinformation posts shared. Tests 1D, 2D, 3D, 4D, 5D, 6D, 7D, 8D, and 9D are one-sided t-tests. Tests 10D, 11D and 12D are two-sided t-tests.

Group 1	Group 2	N. Obs. 1	N. Obs. 2	Mean 1	Mean 2	Diff. in Means	95% CI	Unadj p-Value	Holm Adj p-Value
All	No-course	5,263	1,626	-0.1814 (0.0051)	-0.0269 (0.0084)	-0.1545 (0.0099)	$(-\infty, -0.1383)$	0.0000	0.0000
All	Facts	5,263	1,789	-0.1814 (0.0051)	-0.0824 (0.0083)	-0.0989 (0.0097)	$(-\infty, -0.0829)$	0.0000	0.0000
Info	No-course	1,689	1,626	-0.1504 (0.0089)	-0.0269 (0.0084)	-0.1235 (0.0123)	$(-\infty, -0.1034)$	0.0000	0.0000
Emotions	No-course	1,836	1,626	-0.2039 (0.0089)	-0.0269 (0.0084)	-0.1770 (0.0122)	$(-\infty, -0.1569)$	0.0000	0.0000
Combo	No-course	1,738	1,626	-0.1878 (0.0089)	-0.0269 (0.0084)	-0.1609 (0.0123)	$(-\infty, -0.1407)$	0.0000	0.0000
Info	Facts	1,689	1,789	-0.1504 (0.0089)	-0.0824 (0.0083)	-0.0679 (0.0121)	$(-\infty, -0.048)$	0.0000	0.0000
Emotions	Facts	1,836	1,789	-0.2039 (0.0089)	-0.0824 (0.0083)	-0.1214 (0.0121)	$(-\infty, -0.1015)$	0.0000	0.0000
Combo	Facts	1,738	1,789	-0.1878 (0.0089)	-0.0824 (0.0083)	-0.1053 (0.0122)	$(-\infty, -0.0853)$	0.0000	0.0000
Facts	No-course	1,789	1,626	-0.0824 (0.0083)	-0.0269 (0.0084)	-0.0556 (0.0118)	$(-\infty, -0.0361)$	0.0000	0.0000
Info	Emotions	1,689	1,836	-0.1504 (0.0089)	-0.2039 (0.0089)	0.0535 (0.0126)	$(0.0289, 0.0781)$	0.0000	0.0001
Info	Combo	1,689	1,738	-0.1504 (0.0089)	-0.1878 (0.0089)	0.0374 (0.0126)	$(0.0127, 0.0621)$	0.0030	0.0060
Emotions	Combo	1,836	1,738	-0.2039 (0.0089)	-0.1878 (0.0089)	-0.0161 (0.0126)	$(-0.0408, 0.0085)$	0.1998	0.1998

Table B5: Average Treatment Effects - Non-Corresponding Misinformation Sharing

Notes: The sample is composed of the 8,768 participants who completed the post-survey, pooling participants in the Accuracy Inter and Accuracy After treatments. Each row displays the results of t-tests. For the assignment groups specified in the first column of the table, the t-tests compare the average post-pre difference in the proportion of misinformation posts shared, excluding those misinformation posts in the post-survey that are counterparts of non-misinformation posts in the pre-survey. Tests 1N, 2N, 3N, 4N, 5N, 6N, 7N, 8N, and 9N are one-sided t-tests. Tests 10N, 11N and 12N are two-sided t-tests.

Group 1	Group 2	N. Obs. 1	N. Obs. 2	Mean 1	Mean 2	Diff. in Means	95% CI	Unadj p-Value	Holm Adj p-Value
All	No-course	2,830	862	0.2012 (0.0070)	0.2952 (0.0145)	-0.0940 (0.0161)	$(-\infty, -0.0675)$	0.0000	0.0000
All	Facts	2,830	1,007	0.2012 (0.0070)	0.2337 (0.0124)	-0.0325 (0.0143)	$(-\infty, -0.009)$	0.0114	0.0800
Info	No-course	889	862	0.2135 (0.0129)	0.2952 (0.0145)	-0.0817 (0.0194)	$(-\infty, -0.0497)$	0.0000	0.0001
Emotions	No-course	996	862	0.1954 (0.0116)	0.2952 (0.0145)	-0.0998 (0.0185)	$(-\infty, -0.0693)$	0.0000	0.0000
Combo	No-course	945	862	0.1958 (0.0119)	0.2952 (0.0145)	-0.0995 (0.0187)	$(-\infty, -0.0686)$	0.0000	0.0000
Info	Facts	889	1,007	0.2135 (0.0129)	0.2337 (0.0124)	-0.0202 (0.0179)	$(-\infty, 0.0093)$	0.1305	0.5220
Emotions	Facts	996	1,007	0.1954 (0.0116)	0.2337 (0.0124)	-0.0382 (0.0170)	$(-\infty, -0.0103)$	0.0121	0.0800
Combo	Facts	945	1,007	0.1958 (0.0119)	0.2337 (0.0124)	-0.0379 (0.0172)	$(-\infty, -0.0096)$	0.0138	0.0800
Facts	No-course	1,007	862	0.2337 (0.0124)	0.2952 (0.0145)	-0.0615 (0.0191)	$(-\infty, -0.0301)$	0.0006	0.0051
Info	Emotions	889	996	0.2135 (0.0129)	0.1954 (0.0116)	0.0181 (0.0173)	$(-0.0159, 0.0521)$	0.2971	0.8913
Info	Combo	889	945	0.2135 (0.0129)	0.1958 (0.0119)	0.0178 (0.0176)	$(-0.0167, 0.0522)$	0.3121	0.8913
Emotions	Combo	996	945	0.1954 (0.0116)	0.1958 (0.0119)	-0.0003 (0.0166)	$(-0.0328, 0.0322)$	0.9847	0.9847

Table B6: Average Treatment Effects - Opposite Outcome

Notes: The sample is composed of the 4,699 participants who completed the post-survey and shared at most two non-misinformation posts in the pre-survey, pooling participants in the Accuracy Inter and Accuracy After treatments.. Each row displays the results of t-tests. For the assignment groups specified in the first column of the table, the t-tests compare the ratio of the number of misinformation posts the participant shares in the post-survey that correspond to the non-misinformation counterpart (non-misinformation post) they did not share in the pre-survey to the number of those corresponding non-misinformation posts they did not share in the pre-survey. Tests 1O, 2O, 3O, 4O, 5O, 6O, 7O, 8O, and 9O are one-sided t-tests. Tests 10O, 11O and 12O are two-sided t-tests.

B.6 Alternative Outcomes Family 3 Tests

Group	N Obs.	Mean Pre	Mean Post	Mean Diff.	95% CI	Unadj p-Value	Holm Adj p-Value
Accuracy After	2,654	-1.211 (0.031)	-0.641 (0.031)	0.4977 (0.0385)	(0.4377, ∞)	0.0000	0.0000
Accuracy Inter	2,609	-1.138 (0.031)	-0.686 (0.029)	0.5255 (0.0365)	(0.4621, ∞)	0.0000	0.0000
Diff in Mean Diff	5,263	-	-	-0.0277 (0.0531)	(-0.1318, 0.0763)	0.6012	0.6012

Table B7: Effect of courses and accuracy nudge - Sharing Discernment Score

Notes: The sample is composed of the 5,263 participants in the Reasoning, Emotions, and Combo assignment groups who completed the post-survey. The first two rows display the results of paired one-sided t-tests that compare the post-pre difference in the discernment score of shared posts for the Accuracy After and Accuracy Inter groups, respectively. The last row displays the results of the test of the difference-in-mean-differences coefficient being equal to 0.

Group	N Obs.	Mean Pre	Mean Post	Mean Diff.	95% CI	Unadj p-Value	Holm Adj p-Value
Accuracy After	2,654	1.700 (0.135)	3.146 (0.128)	1.446 (0.1607)	(1.1814, ∞)	0.0000	0.0000
Accuracy Inter	2,609	2.411 (0.137)	3.892 (0.128)	1.4810 (0.1708)	(1.1998, ∞)	0.0000	0.0000
Diff in Mean Diff	5,263	-	-	0.0350 (0.2345)	(-0.4248, 0.4948)	0.8813	0.8813

Table B8: Effect of courses and accuracy nudge - Accuracy Discernment Score

Notes: The sample is composed of the 5,263 participants in the Reasoning, Emotions, and Combo assignment groups who completed the post-survey. The first two rows display the results of paired one-sided t-tests that compare the post-pre difference in the accuracy discernment score of shared posts for the Accuracy After and Accuracy Inter groups, respectively. The last row displays the results of the test of the difference-in-mean-differences coefficient being equal to 0.

B.7 Follow-up Results, by Prime Status




Figure B6: Follow-up, Primed participants




Figure B7: Follow-up, Non-primed participants

C Robustness to sample who correctly answered all attention check questions




Figure B1: Misinformation Sharing, by Intervention Assignment Group - Followup - Attention Check Sample

Notes: Sample includes 2,738 participants who completed the post-survey passing attention checks in the pre- and post- survey, and shared at least one non-misinformation post in the pre-survey. Each bar displays the Sharing Rate for misinformation posts, as defined in Equation 1, by participants in their respective intervention assignment group, pooling participants in the Accuracy Inter and Accuracy After groups. Above each bar, the standard error is shown in parentheses below the Sharing Rate. The thin black bars represent 95% confidence intervals. Differences in Sharing Rates are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

	No-course baseline	Facts baseline	Reasoning course	Emotions course	Combo course	Treatment courses	Baselines	All
Number of Observations								
Accuracy Inter	255	292	273	299	300	872	547	1419
Accuracy After	299	373	340	384	357	1081	672	1753
All	554	665	613	683	657	1953	1219	3172
Pre Non-misinfo Posts								
Accuracy Inter	0.638 (0.021)	0.605 (0.018)	0.648 (0.019)	0.615 (0.018)	0.594 (0.02)	0.618 (0.011)	0.62 (0.014)	0.619 (0.009)
Accuracy After	0.735 (0.02)	0.71 (0.019)	0.722 (0.021)	0.717 (0.02)	0.664 (0.021)	0.7 (0.012)	0.722 (0.014)	0.708 (0.009)
All	0.682 (0.015)	0.651 (0.013)	0.681 (0.014)	0.659 (0.014)	0.626 (0.014)	0.655 (0.008)	0.665 (0.01)	0.659 (0.006)
Pre Misinfo Posts								
Accuracy Inter	0.49 (0.02)	0.449 (0.017)	0.48 (0.019)	0.467 (0.017)	0.458 (0.018)	0.468 (0.01)	0.467 (0.013)	0.468 (0.008)
Accuracy After	0.565 (0.019)	0.567 (0.018)	0.548 (0.021)	0.534 (0.019)	0.509 (0.019)	0.53 (0.011)	0.566 (0.013)	0.544 (0.009)
All	0.524 (0.014)	0.501 (0.013)	0.51 (0.014)	0.497 (0.013)	0.482 (0.013)	0.496 (0.008)	0.511 (0.009)	0.502 (0.006)
Post Non-misinfo Posts								
Accuracy Inter	0.616 (0.022)	0.503 (0.02)	0.522 (0.02)	0.456 (0.019)	0.474 (0.02)	0.483 (0.011)	0.554 (0.015)	0.51 (0.009)
Accuracy After	0.661 (0.023)	0.622 (0.022)	0.554 (0.023)	0.548 (0.022)	0.512 (0.022)	0.538 (0.013)	0.64 (0.016)	0.577 (0.01)
All	0.637 (0.016)	0.555 (0.015)	0.536 (0.015)	0.496 (0.014)	0.492 (0.015)	0.507 (0.009)	0.593 (0.011)	0.54 (0.007)
Post Misinfo Posts								
Accuracy Inter	0.465 (0.019)	0.366 (0.018)	0.369 (0.018)	0.283 (0.015)	0.314 (0.017)	0.32 (0.01)	0.41 (0.013)	0.354 (0.008)
Accuracy After	0.523 (0.021)	0.478 (0.02)	0.421 (0.021)	0.365 (0.019)	0.381 (0.019)	0.388 (0.011)	0.499 (0.015)	0.431 (0.009)
All	0.492 (0.014)	0.415 (0.014)	0.392 (0.014)	0.319 (0.012)	0.344 (0.013)	0.35 (0.007)	0.45 (0.01)	0.389 (0.006)
Primary Outcome: Sharing Rate								
Accuracy Inter	0.609 (0.025)	0.499 (0.024)	0.508 (0.023)	0.384 (0.022)	0.435 (0.024)	0.44 (0.013)	0.549 (0.017)	0.482 (0.011)
Accuracy After	0.624 (0.025)	0.575 (0.026)	0.553 (0.026)	0.433 (0.024)	0.486 (0.026)	0.489 (0.015)	0.598 (0.018)	0.532 (0.012)
All	0.617 (0.018)	0.534 (0.017)	0.528 (0.017)	0.406 (0.016)	0.459 (0.018)	0.463 (0.01)	0.572 (0.012)	0.505 (0.008)
Pre Sharing Discernment								
Accuracy Inter	1.963 (0.419)	3.129 (0.334)	3.018 (0.38)	2.865 (0.322)	2.406 (0.351)	2.761 (0.202)	2.61 (0.264)	2.703 (0.16)
Accuracy After	1.996 (0.414)	1.432 (0.378)	1.571 (0.397)	1.749 (0.392)	2.34 (0.368)	1.897 (0.223)	1.695 (0.279)	1.819 (0.174)
All	1.978 (0.295)	2.383 (0.252)	2.374 (0.276)	2.376 (0.25)	2.376 (0.254)	2.375 (0.15)	2.199 (0.192)	2.308 (0.118)
Post Sharing Discernment								
Accuracy Inter	1.816 (0.367)	3.568 (0.336)	3.324 (0.358)	4.495 (0.31)	4.255 (0.343)	4.047 (0.195)	2.789 (0.25)	3.565 (0.154)
Accuracy After	1.557 (0.424)	2.048 (0.338)	2.802 (0.377)	3.261 (0.343)	2.627 (0.363)	2.899 (0.208)	1.819 (0.268)	2.483 (0.165)
All	1.697 (0.278)	2.901 (0.242)	3.091 (0.26)	3.955 (0.231)	3.511 (0.251)	3.535 (0.143)	2.354 (0.183)	3.081 (0.113)
Pre Accuracy Score Discernment								
Accuracy Inter	1.963 (0.419)	3.129 (0.334)	3.018 (0.38)	2.865 (0.322)	2.406 (0.351)	2.761 (0.202)	2.61 (0.264)	2.703 (0.16)
Accuracy After	1.996 (0.414)	1.432 (0.378)	1.571 (0.397)	1.749 (0.392)	2.34 (0.368)	1.897 (0.223)	1.695 (0.279)	1.819 (0.174)
All	1.978 (0.295)	2.383 (0.252)	2.374 (0.276)	2.376 (0.25)	2.376 (0.254)	2.375 (0.15)	2.199 (0.192)	2.308 (0.118)
Post Accuracy Score Discernment								
Accuracy Inter	1.816 (0.367)	3.568 (0.336)	3.324 (0.358)	4.495 (0.31)	4.255 (0.343)	4.047 (0.195)	2.789 (0.25)	3.565 (0.154)
Accuracy After	1.557 (0.424)	2.048 (0.338)	2.802 (0.377)	3.261 (0.343)	2.627 (0.363)	2.899 (0.208)	1.819 (0.268)	2.483 (0.165)
All	1.697 (0.278)	2.901 (0.242)	3.091 (0.26)	3.955 (0.231)	3.511 (0.251)	3.535 (0.143)	2.354 (0.183)	3.081 (0.113)

Table B1: Summary of Outcomes by Text Message Course Intervention and Accuracy Nudge Assignment - Attention Check Sample

Notes: Sample includes 3,172 participants who completed the post-survey, except for the “Primary Outcome: Sharing Rate” rows that exclude 434 participants who did not share at least one non-misinformation post in the pre-survey. The first two rows display the number of observations in each assignment group. The other rows display averages by assignment group, with standard errors in parentheses.




Figure B2: Misinformation Sharing, by Intervention Assignment Group - Attention Check Sample

Notes: Sample includes 1,796 participants who completed the follow-up survey passing all the attention checks and shared at least one non-misinformation post in the post-survey. Each bar displays the Sharing Rate for misinformation posts, as defined in Equation 1, changing the *post* and *pre* superscripts to *followup* and *post*, respectively, by participants in their respective intervention assignment group, pooling participants in the Accuracy Inter and Accuracy After groups. Above each bar, the standard error is shown in parentheses below the Sharing Rate. The thin black bars represent 95% confidence intervals. Differences in the average Sharing Rates are shown above lines connecting the two relevant intervention assignment groups, with standard errors in parentheses.

D Robustness to controlling for covariates using Augmented Inverse Probability Weighting (AIPW)

Group 1	Group 2	N. Obs. 1	N. Obs. 2	Estimates	95% CI	Unadj p-Value	Holm Adj p-Value
All	No-course	4,646	1,454	-0.1658 (0.0435)	$(-\infty, -0.0942)$	0.0001	0.0003
All	Facts	4,646	1,584	-0.1226 (0.0436)	$(-\infty, -0.0508)$	0.0025	0.0074
Reasoning	No-course	1,502	1,454	-0.1417 (0.0149)	$(-\infty, -0.1172)$	0.0000	0.0000
Emotions	No-course	1,617	1,454	-0.2080 (0.0147)	$(-\infty, -0.1838)$	0.0000	0.0000
Combo	No-course	1,527	1,454	-0.1797 (0.0151)	$(-\infty, -0.1548)$	0.0000	0.0000
Reasoning	Facts	1,502	1,584	-0.0736 (0.0150)	$(-\infty, -0.0489)$	0.0000	0.0000
Emotions	Facts	1,617	1,584	-0.1402 (0.0148)	$(-\infty, -0.1158)$	0.0000	0.0000
Combo	Facts	1,527	1,584	-0.1103 (0.0151)	$(-\infty, -0.0855)$	0.0000	0.0000
Facts	No-course	1,584	1,454	-0.0699 (0.0151)	$(-\infty, -0.0451)$	0.0000	0.0000
Reasoning	Emotions	1,502	1,617	0.0656 (0.0147)	$(0.0368, 0.0944)$	0.0000	0.0000
Reasoning	Combo	1,502	1,527	0.0345 (0.0152)	$(0.0047, 0.0643)$	0.0231	0.0463
Emotions	Combo	1,617	1,527	-0.0242 (0.0148)	$(-0.0532, 0.0047)$	0.1013	0.1013

Table C1: Average Treatment Effects - AIPW

Notes: The sample is composed of the 7,774 participants who completed the post-survey and shared at least one non-misinformation post in the pre-survey, pooling participants in the Accuracy Inter and Accuracy After treatments. Each row displays the AIPW treatment effect estimates for the primary outcome in Equation 1. The first nine rows show one-sided tests, while the last three rows comparing treatment courses to each other shows two-sided tests.

Group 1	Group 2	N. Obs. 1	N. Obs. 2	Estimates	95% CI	Unadj p-Value	Holm Adj p-Value
All	Facts	2,554	857	-0.1103 (0.0607)	$(-\infty, -0.0105)$	0.0345	0.1380
Reasoning	Facts	857	857	-0.0726 (0.0207)	$(-\infty, -0.0386)$	0.0002	0.0013
Emotions	Facts	865	857	-0.1286 (0.0208)	$(-\infty, -0.0944)$	0.0000	0.0000
Combo	Facts	832	857	-0.0814 (0.0208)	$(-\infty, -0.0471)$	0.0000	0.0003
Facts	No-course	857	838	0.0204 (0.0214)	$(-0.0216, 0.0623)$	0.3414	0.6829
Reasoning	Emotions	857	865	0.0501 (0.0204)	$(0.0101, 0.0901)$	0.0140	0.0701
Reasoning	Combo	857	832	0.0105 (0.0206)	$(-0.0299, 0.0509)$	0.6119	0.6829
Emotions	Combo	865	832	-0.0357 (0.0207)	$(-0.0762, 0.0048)$	0.0838	0.2515

Table C2: Average Treatment Effects - AIPW Follow-up

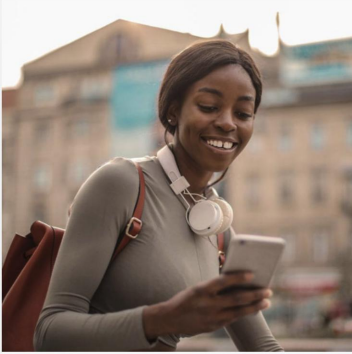
Notes: The sample is composed of the 4,249 participants who completed the follow-up survey and shared at least one non-misinformation post in the post-survey, pooling participants in the Accuracy Inter and Accuracy After treatments. Each row displays the Augmented Inverse Propensity Weighting treatment effect estimates on the ratio of the number of misinformation posts the participant shares in the follow-up survey that correspond to the non-misinformation counterpart (non-misinformation post) they shared in the post-survey to the number of those corresponding non-misinformation posts they did share in the post-survey, for the assignment groups specified in the first column of the table. Tests 2AF, 6AF, 7AF, and 8AF are one-sided. Tests 9AF, 10AF, 11AF, and 12AF are two-sided.

Online Appendix A: Facebook Recruitment Ads Examples

Social Impact Research Lab Sponsored · 🌐 ...

Let's make social media better together!

Earn 🕒 AIRTIME 🕒 by signing up for our FREE, five-day text message course and completing short course surveys.



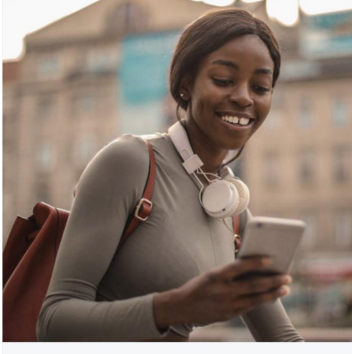
stanforduniversity.qualtrics.com
Start now [Sign up](#)

👍 Joy Muthoni Muir... 11 Comments

👍 Like 💬 Comment ➦ Share

Social Impact Research Lab Sponsored · 🌐 ...

Earn 🕒 AIRTIME 🕒 by signing up for our FREE, five-day text message course and completing short course surveys.



stanforduniversity.qualtrics.com
Start now [Sign up](#)

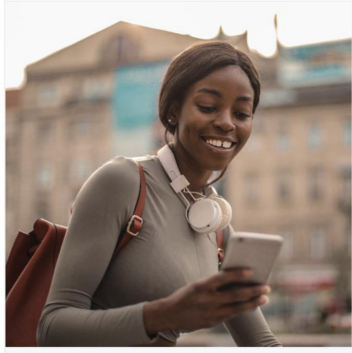
👍 Thomas Arisa and 2... 2 Comments 1 Share

👍 Like 💬 Comment ➦ Share

Social Impact Research Lab Sponsored · 🌐 ...

Do you want to protect your family and friends from misinformation online?

Earn 🕒 AIRTIME 🕒 by signing up for our FREE, five-day text message course and completing short course surveys.



stanforduniversity.qualtrics.com
Start now [Sign up](#)

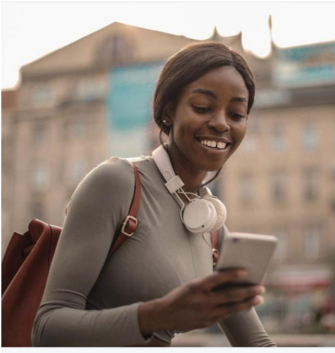
👍 Magdaline Adero and 9... 4 Comments

👍 Like 💬 Comment ➦ Share

Social Impact Research Lab Sponsored · 🌐 ...

Not sure how to protect yourself from misinformation?

Earn 🕒 AIRTIME 🕒 by signing up for our FREE, five-day text message course and completing short course surveys.



stanforduniversity.qualtrics.com
Start now [Sign up](#)

👍 Eli Muthaura and 25 others

👍 Like 💬 Comment ➦ Share

Online Appendix B: Text Message Courses

Emotions Course

DAY 1: Why people share misleading content

Let's start by talking about some definitions. Do you know the difference between disinformation and misinformation?

The difference is INTENT.

👉 DISINFORMATION is false or misleading information that is deliberately created and shared by people who know that it is false.

👉 MISINFORMATION is false or misleading information that people share without necessarily knowing that it is false.

Sometimes people share misleading information because they are trying to help, but it's important to remember the harm that can be caused whenever we share anything that's not 100% true.

? Let's test what you just learned. Pretend you have just seen a post on Facebook from your uncle sharing what he believes is an alternative treatment for Covid-19. You know that this treatment has been disproven. Is this an example of misinformation or disinformation?

A. Misinformation

✅ Correct! This is an example of MISINFORMATION. More likely than not, if your uncle shares an alternative treatment for Covid-19, he is probably just trying to help.

B. Disinformation

❌ Almost! This is an example of misinformation. More likely than not, if your uncle shares an alternative treatment for Covid-19, he is probably just trying to help. Disinformation is false or misleading information that is deliberately shared by people who know that it is false.

There are a lot of reasons why people are trying to deliberately mislead you:

- 💰 To make money
- 🏆 To build reputation
- 😈 To cause trouble
- 📊 Political gain

➡ Next time you see something misleading online, ask yourself if it's MISinformation or DISinformation, and think about why that person might have shared it!

Tomorrow, we'll talk about how our brains process misinformation. We'll also learn about how we can watch out for it.

Goodbye for now! 🙌

DAY 2: How to protect ourselves

Let's talk about our brains. 🧠

🧑 Our brains help us make decisions all day long. But there are also some things they do that make us more susceptible to misinformation.

1 EMOTIONAL TRIGGERS

— Ever heard the phrase “fight or flight”? When we're scared, we do whatever we need to protect ourselves and others. 🗨️ That includes consuming and sharing information we believe will keep us safe but which might not actually be true.

— The same is true for information that makes us angry, 😡 scared 😱 or sad 😞. When we have strong emotional reactions, we're less likely to stop, think and analyze what we're seeing.

2 MENTAL SHORTCUTS

— We also use a lot of mental shortcuts to help us make snap decisions. 🏃 For example, if a friend posts something misleading on social media, we might be more likely to share it because we trust that person.

— Here's the problem: People who spread disinformation often rely on these emotional triggers and mental shortcuts to manipulate us. 😞 Your friend might not know they are being misled, but you can help stop the spread by not sharing. 🚫

? How about you? If a family member or a friend posts something online, do you stop to check whether it's misinformation?

A. I always trust my family and friends!

That's normal! But remember: Everyone makes mistakes. It's always a good idea to verify whether something is true before resharing it. ✅

B. I always check whether something is misinformation.

That's great! Keep it up. 🎯

➡ One of the best ways to protect ourselves from misinformation is to recognize when we're having a strong emotional response. Try following these steps:

1 🛑 STOP: If you read or see something online that makes you feel strongly, pause before sharing it.

2 🤔 QUESTION: Is the post's creator trying to manipulate how you think about something? Is the information misleading? If there is any doubt in your mind, do NOT share it.

Remember that by sharing misleading information, you could cause real harm. Just like you trust many of your friends and family, they trust you too! 💪

See you tomorrow! 🌟

DAY 3: Fear



It's time for our next lesson! Over the next three days, we're going to show you a few examples of how misinformation targets your emotions.

Today let's talk about FEAR.

👹 Life can seem pretty scary right. Feeling scared or uncertain is rational, but it's important to understand that when we are scared we are also more vulnerable to being manipulated. 🚫

Take a look at the image above.

Remember the two steps?

1 🛑 STOP: The image suggests that the ingredients in the medicine are unnatural or unsafe — that sounds scary! But let's take a second to think about it.

2 🤔 QUESTION: Is the post's creator trying to manipulate how you think about the medicine?

? What motive might someone have for wanting you to believe the ingredients of another product are safer?

A. To sell their product.

- B. Because they genuinely think the other product is dangerous.
- C. All of the above.

The correct answer is all of the above! ✅ In this example, the person who made this image is clearly trying to make a profit off the alternative product. 📊 But, of course, the person also might genuinely think the other product is dangerous.

➔ In either case, notice how the image makes you FEEL. Health information can be difficult to understand because of unfamiliar terms, and unfamiliarity often breeds fear. 🧑

But remember: Fear makes you more vulnerable to being manipulated, and people who spread disinformation often use this to their advantage.

Stop and seek out all the information you might need to make informed decisions. And if in doubt, do NOT share. 🙅

That's all for now! See you tomorrow. 🙌

Day 4: Anger

**DON'T LET THE GOVERNMENT MAKE
YOUR CHILD AN EXPERIMENT**



Welcome back! Today let's talk about ANGER.

Our beliefs are part of our identity, and when these beliefs are attacked we often become angry and defensive. 😡

👉 Anger is also a mobilizing force. Research shows that misinformation travels much faster on social media if it provokes anger.

Some platforms have even prioritized showing users content that elicits angry reactions! This makes it much harder for us to avoid content that is polarizing. 🙄

? Take a look at the image above. Would you share this on your social media feed?

- A. Yes.
- B. No.

Before sharing, let's stop and question all the ways this post tries to make us feel anger:


- ➡ The picture shows a crying child. People are naturally protective of their children.
- ➡ The caption implies that the vaccination being administered is experimental.
- ➡ The caption also implies that the "government" does not care about children's health or safety.

Instead of providing useful information to help you understand whether the vaccine depicted is safe and why, this post is relying on emotional triggers to make you angry.

🚨 Remember: Before sharing something that makes you angry, STOP and QUESTION the information. Don't share anything if you're not 100% sure it is true.

See you for our final lesson tomorrow! 🙌

Day 5: Superiority



Hi again! Today is our last day, and we're going to talk about SUPERIORITY.

There may come a time when you believe that you are privy to information that no one else has access to — that you are part of a select few who sees the “truth.” 👁

It's important to notice when this happens! Here's why:

🧠 Our brains look for clues to quickly work out whether something is credible.

👉 Having access to insider or expert knowledge about a particular issue, or the illusion of having access to this information, is one example of this.

These feelings of “knowing more” than other people might make us feel good, but they can also make us less critical and more prone to believing misinformation. ⚠

➡ Here's one tip. Watch out for health information on social media that begins with language like this:

—1 “My friend is a nurse and she said...”

—2 “My brother works for the government and has inside knowledge. He just told me that...”

🙄 The next time you're presented with "insider information" that someone claims is from an expert, try to understand WHY that person is an expert and see whether there are other people in their field who agree with them.

-

! IMPORTANT ! In order to complete the course and receive free airtime, take a second to fill out our final survey. Are you ready?

A. Yes

Great! Here is the link:

B. No

That's OK! We will remind you again tomorrow. Here's the link if you change your mind.

DAY 6: End-of-course survey

! IMPORTANT ! Did you fill out the end-of-course survey yesterday? In order to receive KSH 500 in AIRTIME, you MUST fill out the survey:

Completing the survey also helps us improve our course and research. 🙏

Congratulations! In this course, you've learned about how misinformation spreads, how it can trick your brain and how to avoid it. 🎉

The next time you're scrolling through your social media feeds and something jumps out at you, stop, think and remember some of the tips you've learned. Doing so could protect you and the ones you love from possible manipulation. 💪

Goodbye! 🙌

Info Course

DAY 1: Why people share misleading content

Let's start by talking about some definitions. Do you know the difference between disinformation and misinformation?

The difference is INTENT.

👉 **DISINFORMATION** is false or misleading information that is deliberately created and shared by people who know that it is false.

👉 **MISINFORMATION** is false or misleading information that people share without necessarily knowing that it is false.

Sometimes people share misleading information because they are trying to help, but it's important to remember the harm that can be caused whenever we share anything that's not 100% true.

? Let's test what you just learned. Pretend you have just seen a post on Facebook from your uncle sharing what he believes is an alternative treatment for Covid-19. You know that this treatment has been disproven. Is this an example of misinformation or disinformation?

A. Misinformation

✅ Correct! This is an example of MISINFORMATION. More likely than not, if your uncle shares an alternative treatment for Covid-19, he is probably just trying to help.

B. Disinformation

❌ Almost! This is an example of misinformation. More likely than not, if your uncle shares an alternative treatment for Covid-19, he is probably just trying to help. Disinformation is false or misleading information that is deliberately shared by people who know that it is false.

There are a lot of reasons why people are trying to deliberately mislead you:


- 💰 To make money
- 🏆 To build reputation
- 😡 To cause trouble
- 📊 Political gain

➔ Next time you see something misleading online, ask yourself if it's MISinformation or DISinformation, and think about why that person might have shared it!

☀️ Tomorrow, we'll talk a little bit about how our brains process misinformation. We'll also learn about how we can watch out for it.

Goodbye for now! 🙌

DAY 2: Misleading graphs



🙌 Welcome back! Over the next few days, we're going to talk about clues and signs to watch for, so that you can get better at spotting mis- and disinformation. 🧑

As the old saying goes: "A picture is worth a thousand words." 🖼️

⚠️ This is especially true when it comes to misinformation.

📊 Graphs and diagrams often make information look more official. The danger is that they can easily be misleading, especially if they use cherry-picked statistics and don't provide all of the data and context.

Misleading graphs or diagrams may be spread intentionally. But even accurate diagrams, when done poorly, can be misinterpreted. 🧐

When you see a graph or diagram on social media, be sure to look it over carefully.

👉 Check the source, where the data comes from, as well as the context in which it was originally shared.

? Look at the chart above. Can you tell what's misleading about it?

- A. There has been no decline in flu vaccine uptake among university students
- B. The numbers have been fabricated
- C. The chart is set up in a way that exaggerates the decline

Automated response:

The correct answer is C! Look carefully at the vertical axis. It starts at 45,000 instead of 0. If the axis started at zero, the decline in vaccine uptake would look less dramatic! This often happens with infographics. The data being displayed might be correct, but the *way* it is displayed is misleading.

That's all for now! We'll see you tomorrow, with a new tactic to learn about. 🧠👉

DAY 3: Websites and impostors

📄 Let's talk about ways people can make misleading content look more "official." A WEBSITE is one way to do this.

Websites are easy to make and they provide a quick way to monetize misinformation by connecting content to advertising. 🤑

📄 Websites that post misinformation are often made to look like an established news or health authority site. This is what we call "impostor content." 😈

➡ Impostor content is when somebody makes it look as if their information came from an organization people recognize and trust. Somebody could do this by using another brand's logo or simply having a similar name or visuals.

🔍 When you see an unfamiliar website or social media page that is pushing health information, try to figure out who created it.

- 🤔 What does the “About” page say?
- 🤔 Who contributes the information?
- 🤔 What are their credentials and background?

Here are some red flags:

- 🚩 The website uses content that comes from somewhere else
- 🚩 The website is selling a health product
- 🚩 The website URL has an unusual ending
- 🚩 The contributors claim to be doctors but you can't find their medical credentials

Goodbye for now! 🙋

Day 4: Eyewitness media

🙋 Welcome back!

Today we're going to talk about the power of eyewitness media. 👁

🗣️ Videos and audio that make it look as if someone experienced something first-hand can be VERY convincing.

📺 Throughout the Covid-19 pandemic, people have shared videos they claim to have filmed showing empty hospitals or people having adverse reactions to vaccines.

These videos are often highly emotive and may depict shocking material. ⚠️ But they are also frequently taken out of context! Here are two examples:

- ➡️ The video or picture might be real, but the caption is misleading.
- ➡️ The video or picture might be old, and it's being reshared to represent something happening now. 🔄

It's always a good idea to verify pictures and videos before you reshare them. Here are some questions to ask yourself:

- ✓ Is the person sharing a picture or video the person who took the picture or video?
- ✓ Is it possible that the picture or video came from somewhere else?
- ✓ If so, where did it come from? Can you find other examples of it on the internet?
- ✓ When was the picture or video taken?

—

? Want to know about a tool you can use to look up where a photo came from?

A. Yes! Show me the tool.

This is a free tool that works on both Google Chrome and Firefox. When you install it, it allows you to right-click on any photo and do what's called a reverse image search.

<https://chrome.google.com/webstore/detail/reveye-reverse-image-sear/keaacjcjhehbbapnphnmpiklalfhelgf?hl=en>

B. That's OK! I've learned enough for one day.

OK, great!

Bye for now! 🙌

Day 5: The kernel of truth

CONCERT IN NAIROBI

Are they crazy? No one in this photo has a mask on!



Today we're going to talk about one of the most important concepts: CONTEXT.

🤖 The reality is that most misinformation isn't entirely false. A lot of misinformation has a small kernel of truth that has been presented in a false or misleading way. 🔍

!! 🗣️ This is important, because it's often the small kernel of truth that makes the content more believable.

💬 Examples of this kind of misinformation include content taken out of context, such as a quote that is cut and reframed to change the meaning. ✂️

Because the content contains some degree of truth, it is more complicated to debunk.

👉 Remember: Context is important. Ask yourself not only whether particular facts are true, but also whether they are being presented accurately.

—

? Say you come across an image like the one above. What do you think is the most likely explanation for why it's misleading?

- A. It has been Photoshopped
- B. It is out of context

Sometimes images are Photoshopped, but most of the misinformation you will come across is not outright fabricated, it's just out of context. This could be an old photo of a maskless crowd from 2015, depicted as if it had been taken during the pandemic.

! IMPORTANT ! In order to receive airtime, take a second to fill out our final survey. Type 'A' if you're ready for the link.

Great! Please take our end-of-course survey here:

DAY 6: End-of-course survey

! IMPORTANT ! Did you fill out the end-of-course survey yesterday? In order to receive airtime, you MUST fill out the survey.

🙏 Completing the survey also helps us improve our course and research.

Congratulations! In this course, you've learned about how misinformation spreads, how it can trick your brain and how to avoid it. 🎉

The next time you're scrolling through your social media feeds and something jumps out at you, stop, think and remember some of the tips you've learned. Doing so could protect you and the ones you love from possible manipulation.

Goodbye! 👋

Combo Course

DAY 1: Why people share misleading content

Let's start by talking about some definitions. Do you know the difference between disinformation and misinformation?

The difference is INTENT.

👉 **DISINFORMATION** is false or misleading information that is deliberately created and shared by people who know that it is false.

👉 **MISINFORMATION** is false or misleading information that people share without necessarily knowing that it is false.

Sometimes people share misleading information because they are trying to help, but it's important to remember the harm that can be caused whenever we share anything that's not 100% true.

? Let's test what you just learned. Pretend you have just seen a post on Facebook from your uncle sharing what he believes is an alternative treatment for Covid-19. You know that this treatment has been disproven. Is this an example of misinformation or disinformation?

A. Misinformation

✅ Correct! This is an example of MISINFORMATION. More likely than not, if your uncle shares an alternative treatment for Covid-19, he is probably just trying to help.

B. Disinformation

❌ Almost! This is an example of misinformation. More likely than not, if your uncle shares an alternative treatment for Covid-19, he is probably just trying to help. Disinformation is false or misleading information that is deliberately shared by people who know that it is false.

There are a lot of reasons why people are trying to deliberately mislead you:

💰 To make money

🗣️ To build reputation

😡 To cause trouble

📊 Political gain

➡️ Next time you see something misleading online, ask yourself if it's MISinformation or DISinformation, and think about why that person might have shared it!

🌟 Tomorrow, we'll talk about how our brains process misinformation. We'll also learn about how we can watch out for misinformation.

👋 Goodbye for now!

DAY 2: How to protect ourselves

Let's talk about our brains. 🧠

🧑 Our brains help us make decisions all day long. But there are also some things they do that make us more susceptible to misinformation.

1 EMOTIONAL TRIGGERS

— Ever heard the phrase “fight or flight”? When we're scared, we do whatever we need to protect ourselves and others. 🗣️ That includes consuming and sharing information we believe will keep us safe but which might not actually be true.

— The same is true for information that makes us angry, 😡 scared 😱 or sad 😞. When we have strong emotional reactions, we're less likely to stop, think and analyze what we're seeing.

2 MENTAL SHORTCUTS

— We also use a lot of mental shortcuts to help us make snap decisions. 🏃
For example, if a friend posts something misleading on social media, we might be more likely to share it because we trust that person.

— Here's the problem: People who spread disinformation often rely on these emotional triggers and mental shortcuts to manipulate us. 😞 Your friend might not know they are being misled, but you can help stop the spread by not sharing. 🚫

? How about you? If a family member or friend posts something online, do you stop to check whether it's misinformation?

A. I always trust my family and friends!

That's normal! But remember: Everyone makes mistakes. It's always a good idea to verify whether something is true before resharing it. ✅

B. I always check whether something is misinformation.

That's great! Keep it up. 🏆

➔ One of the best ways to protect ourselves from misinformation is to recognize when we're having a strong emotional response. Try following these steps:

1 🙅 STOP: If you read or see something online that makes you feel strongly, pause before sharing it.

2 🤔 QUESTION: Is the post's creator trying to manipulate how you think about something? Is the information misleading? If there is any doubt in your mind, do NOT share it.

Remember that by sharing misleading information, you could cause real harm. Just like you trust many of your friends and family, they trust you too! 🙏

DAY 3: Paying attention to our emotions



Let's talk about some ways that misinformation targets your emotions.

- Fear
- Anger
- Superiority

1 FEAR 🤖

Life can seem scary right now. Many people have been seriously ill in the pandemic, and we're not sure what the future holds.

🔥 Feeling scared or uncertain is rational, but it's important to understand that when we are scared we are also more vulnerable to being manipulated.

2 ANGER 😡

It has also been shown that misinformation spreads faster if it provokes anger.

Some platforms have even prioritized showing users content that receives angry reactions! This makes it much harder for us to avoid content that is polarizing. 🤖

③ SUPERIORITY 😊

There may come a time when you believe that you are privy to information that no one else has access to — that you are one of a select few who see the “truth.” 👁

These feelings of “knowing more” might make us feel good, but they can also make us less critical ⚠

➡ Here’s one tip. Watch out for health information on social media that begins with language like this:

- 1 “My friend is a nurse and she said...”
- 2 “My brother works for the government. He told me...”

? Take a look at the image above. What motive might someone have for making you believe the ingredients of another product are safer?

- A. To sell their product.
- B. To make you scared of the other product.
- C. Because they genuinely think the other product is dangerous.
- D. All of the above.

The correct answer is all of the above! ✅ In this example, the person who made this image is clearly trying to make a profit off their alternative product. 🏠 But, of course, the person also might genuinely think the other product is dangerous!

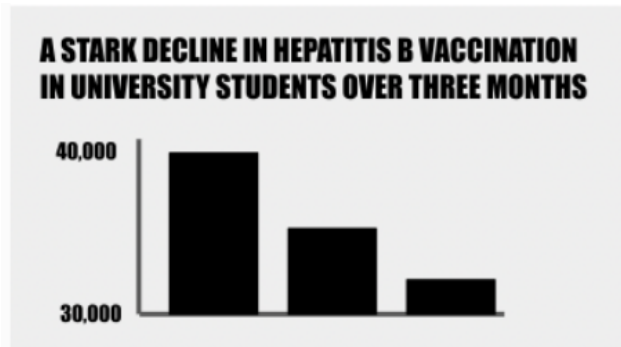
➡ In either case, notice how the image makes you FEEL. Health information can be difficult to understand because of unfamiliar terms, and unfamiliarity often breeds fear. 🧑

But remember: Fear makes you more vulnerable to being manipulated, and people who spread disinformation often use this to their advantage.

Stop and seek out all the information you might need to make informed decisions. And if in doubt, do NOT share. 💪

Bye for now! 🙌

Day 4: Misinformation tactics



Let's talk about some of the things that can make misinformation look so credible.

- Websites and impostors
- Eyewitness media
- Graphs and diagrams

1 WEBSITES AND IMPOSTORS

Websites are easy to make and they provide a quick way to monetize misinformation by connecting content to advertising. 💰

📰 Websites that post misinformation are often made to look like an established news or health authority site. This is what we call "impostor content." 😡

2 "EYEWITNESS" MEDIA

Videos that make it look as if someone experienced something first-hand can be VERY convincing. 📹 Throughout the Covid-19 pandemic, people have shared videos they claim to have filmed showing empty hospitals or people having adverse reactions to vaccines.

These videos often depict shocking material. ⚠️ But they are also frequently taken out of context!

3 GRAPHS AND DIAGRAMS

📊 Graphs and diagrams often make information look more official. The danger is that they can easily be misleading, especially if they use cherry-picked statistics and don't provide all of the data and context.

Even accurate diagrams, when done poorly, can be misinterpreted. 🧑

- ? Take a look at the chart above. Can you tell what's misleading about it?
- A. There has been no decline in flu vaccine uptake among university students
 - B. The numbers have been fabricated
 - C. The chart is set up in a way that exaggerates the decline

Automated response:

The correct answer is C! Look carefully at the vertical axis. It starts at 45,000 instead of 0. If the axis started at zero, the decline in vaccine uptake would look less dramatic! This often happens with infographics. The data being displayed might be correct, but the *way* it is displayed is misleading.

That's all for now! See you tomorrow for our last lesson. 🙌

Day 5: The kernel of truth

CONCERT IN NAIROBI

Are they crazy? No one in this photo has a mask on!



Today we're going to talk about one of the most important concepts: CONTEXT.

🤔 The reality is that most misinformation isn't entirely false. A lot of misinformation has a small kernel of truth that has been presented in a false or misleading way. 🔍

!! 🗣️ This is important, because often the small kernel of truth makes the content more believable.

💬 Examples of this kind of misinformation include content taken out of context, such as a quote that is cut and reframed to change the meaning. ✂️

Because the content contains some degree of truth, it is more complicated to debunk.

➡️ Remember: Context is important. Ask yourself not just whether particular facts are true, but also whether they are being presented accurately.

—

? Say you come across an image like the one above. What do you think is the most likely explanation for why it's misleading?

- A. It has been Photoshopped
- B. It is out of context

Sometimes images are Photoshopped, but most of the misinformation you will come across is not outright fabricated, it's just out of context. This could be an old photo of a maskless crowd from 2015 depicted as if it had been taken during the pandemic.

! IMPORTANT ! In order to receive airtime for taking this course, take a second to fill out our final survey.

DAY 6: End-of-course survey

! IMPORTANT ! Did you fill out the end-of-course survey yesterday? In order to receive KSH 500 in AIRTIME, please take a second to fill out our final survey:

🙌 Completing the survey also helps us improve our course and research.

Congratulations! In this course, you've learned about how misinformation spreads, how it can trick your brain and how to avoid it. 🎉

The next time you're scrolling through your social media feeds and something jumps out at you, stop, think and remember some of the tips you've learned. Doing so could protect you and the ones you love from possible manipulation.

Goodbye! 🙋

Facts Baseline Course

DAY 1:

? Did you know that "lies spread faster than the truth" on social media?

📄 A study from 2018 examined how both true and false news stories spread on Twitter. The authors of the study showed how false news stories traveled farther and faster on the social media platform than verified news stories.

📄 They also showed that of all the different kinds of false news stories, such as false news stories about science or natural disasters, false political news spread the fastest.

📄 Did you also know that being proficient with social media does not necessarily correlate to being skillful in detecting misinformation?

📄 A study from 2021, showed that high school students in the US, despite their frequent use of social media, were largely unable to discern whether a piece of content was true or false.

-

? Let's test what you just learned. Does being proficient on social media correlate to being good at detecting misinformation?

A. Yes

✗ It turns out there is NOT a correlation between being proficient on social media and being skillful in detecting misinformation. A study 📄 from 2021, showed that high school students in the US, despite their frequent use of social media, were largely unable to discern whether a piece of content was true or false.

B. No, not necessarily

✅ That's right. There is NOT a correlation between being proficient on social media and being skillful in detecting misinformation. A study 📄 from 2021, showed that high school students in the US, despite their frequent use of social media, were largely unable to discern whether a piece of content was true or false.

DAY 2:

Many people may believe that misinformation and disinformation are new phenomena that came about with the rise of social media.

!! But misinformation isn't exactly a new thing.

👤 Social media and advanced communication technology may have allowed for misinformation to circulate wider and faster than ever before, but misinformation and disinformation have a long history alongside humankind.

👤 During the Roman Empire, for example, leaders used misinformation to come to power. Nazi Germany was known for its use of propaganda and disinformation. And much of the sensationalist tabloid or "yellow journalism", which effectively treated rumors and falsehoods as facts, was popular in the mid-1900s.

-

Is misinformation a new phenomena?

A. Yes

❌ Misinformation has actually been around for a long time. Yet, social media and advanced communication technology 🧑 may have allowed for misinformation to circulate wider and faster than ever before.

B. No, not necessarily

✅ That's right. Misinformation has a long history alongside humankind. Leaders of the Roman Empire used it as did Nazi Germany. There are many historical examples of misinformation.

DAY 3:

😬 For those of you who think you are unaffected by false information or would know if misinformation was influencing your behavior, guess again!!

📖 Several studies have shown that disinformation, even the slightest amount, can affect your unconscious behavior in very subtle ways.

🔪 Misinformation can also influence your reasoning and decision-making resulting in poor judgment.

🧠 Because of the way our brain operates, even when misinformation we have been exposed to is corrected with factual information, we often continue to be biased towards misinformation. It's what psychologists refer to as the "continued influence effect".

-

? Pretend someone was exposed to a piece of misinformation which they believed. The next day they read a fact-check that debunks the misinformation. What happens to that person?

A. The person is no longer influenced by the piece of misinformation.

✗ Because of the way our brain 🧠 operates, even when misinformation we have been exposed to is corrected with factual information, we often continue to be biased towards misinformation. It's what psychologists refer to as the "continued influence effect".

B. The person continues to be influenced by misinformation.

✓ That's right. Because of the way our brain 🧠 operates, even when misinformation we have been exposed to is corrected with factual information, we often continue to be biased toward misinformation. It's what psychologists refer to as the "continued influence effect".


DAY 4:

👤 In the last lesson, we learned about the "continued influence effect", which shows how we continue believing misinformation even after it has been corrected.

🙄 In today's lesson we will learn about another phenomenon called the "third-person effect", which is particularly relevant to all of us who think that we are immune to misinformation.

👉 The "third-person effect" describes how individuals perceive the effects of misinformation to be much more powerful on others than on themselves. In essence, we think that everyone is susceptible to misinformation except ourselves.




📖 Yet, studies show that people who are the most confident at being able to spot misinformation are the ones most likely to believe it. This is especially true when we are subjected to information overload — when we feel overwhelmed by the amount of information we are exposed to.

 Keep that in mind the next time you are reading the news on social media or have hundreds of tabs open on your computer.


-

? How accurate are humans when it comes to determining the effects of misinformation on ourselves?




A. Super accurate.

 You are super wrong. The “third-person effect” describes how individuals perceive the effects of misinformation to be much more powerful on others  than on themselves . In essence, we think that everyone is susceptible to misinformation except ourselves.


We think that everyone is susceptible to misinformation except ourselves.

Yet, studies  show that people who are the most confident at being able to spot misinformation are the ones most likely to believe it.




B. Pretty accurate.

 Not so much. The “third-person effect” describes how individuals perceive the effects of misinformation to be much more powerful on others  than on themselves . In essence, we think that everyone is susceptible to misinformation except ourselves.


We think that everyone is susceptible to misinformation except ourselves.

Yet, studies  show that people who are the most confident at being able to spot misinformation are the ones most likely to believe it.


C. Not so accurate.


 That's correct. The “third-person effect” describes how individuals perceive the effects of misinformation to be much more powerful on others  than on themselves . In essence, we think that everyone is susceptible to misinformation except ourselves.

We think that everyone is susceptible to misinformation except ourselves.


Yet, studies  show that people who are the most confident at being able to spot misinformation are the ones most likely to believe it.


DAY 5:


 We may think of misinformation as being specific to a particular context or country.

 For example, you might think that the specific rumors and misinformation about a particular medicine or vaccine in one country wouldn't appear in another country.

But the internet and social media are, to a large degree, borderless. Information, including misinformation and disinformation, travels seamlessly between countries and continents.

 Language is one good predictor of where misinformation may spread to. Misinformation in one language, such as English, will more easily make its way to other countries and communities that speak or understand English.

 But language isn't the only predictor. If a piece of misinformation is popular enough, it is often translated word-for-word into other languages, taking on a life of its own.

 Disinformation campaigns are no different. Where bad actors in different countries share similar beliefs or ideologies, they may help each other to influence matters in those countries.

Misinformation is truly a global phenomenon affecting all of us.

-

! IMPORTANT ! In order to complete the course and receive KSH 500 in AIRTIME, take a second to fill out our final survey. Are you ready?

A. Yes.


Great! Here is the link:

B. Not yet.


That's okay! We will remind you again tomorrow. Here's the link if you change your mind.


DAY 6:

! IMPORTANT ! Did you fill out the end-of-course survey yesterday? In order to receive KSH 500 in AIRTIME, please take a second to fill out our final survey.

 Congratulations! In this course, you've learned about how misinformation spreads, how it can trick your brain, and how to avoid it.

The next time you're scrolling through your social media feeds and something jumps out at you, stop, think and remember some of the tips you've learned. Doing so could protect you and the ones you love from possible manipulation.

 If you haven't already, please help us improve our course and research by completing our end-of-course survey:

 If you would like to receive KSH 500 in AIRTIME for completing this course, you **MUST** fill out the survey.

Online Appendix C: Pre-survey Screenshots

D.1 Main Survey

WELCOME!

You are invited to participate in a research study on news. There will be two surveys and a short and free text message course. Your participation will take approximately 10 - 15 minutes for each survey. You will receive the second survey after approximately 5 days after you finish this survey.

You will receive **KSH 500 in AIRTIME** as a thank you for your participation in this research study. You'll be paid within 24 hours of completing the second survey.

Also, make sure you have not completed this survey before. If you complete it more than once, you will not be able to complete the study and receive AIRTIME.

Before we start, please read the consent form and confirm you want to participate in the next screen.



CONSENT FORM

This is a minimal-risk study, and all of your information is confidential. We cannot promise that you will receive any benefits from this study. Your participation is voluntary and you have the right to refuse to participate, skip or refuse to answer any particular questions, or stop at any time.

If you have any questions or concerns, you can contact the Protocol Director, Herman Donner, at inoculationagainstmisinfo@gmail.com / hdonner@stanford.edu or our research review board to speak to someone independent of the research team at irbnonmed@stanford.edu.

If you agree to participate in this research and confirm you are 18 years or older, please click "Yes, I agree".

Yes, I agree

No



You will see 10 social media posts on the following screens. Please read each one of them and answer the question at the end of each screen.

We will include **AT LEAST ONE** *attention check*, which will be in the format of a social media post. When you see the attention check(s), choose the answer as instructed in the post.

If you don't answer ALL the attention checks correctly, you will not receive AIRTIME regardless of whether you completed the survey.

Also, make sure you have not completed this survey before. If you complete it more than once, you will not be able to complete the study and receive AIRTIME.

Please confirm your understanding of the instructions above.

- I acknowledge that I will not receive airtime if I do not pass the attention checks, regardless of whether I complete the survey or not, or if I complete this survey more than once.



Please read the post and answer the question below.




Ali Ouma

8 hrs



This is OUTRAGEOUS! The government is pushing COVID-19 vaccines that don't even protect against deadly COVID-19 infections.



To the best of your knowledge, how accurate is the claim in the above post?

Not at all accurate

Not very accurate

Somewhat accurate

Very accurate



Please read the post and answer the question below.




Ali Ouma

8 hrs



This is OUTRAGEOUS! The government is pushing COVID-19 vaccines that don't even protect against deadly COVID-19 infections.



Would you share this post?

No

Yes





Grace Wangari

5 hrs



Lupita Nyong'o looks gorgeous for 39 years old! Apparently her diet (no sugar, limited carbs and calories, low inflammatory, etc) is the latest in longevity science. Inspiring!



In a previous question, you said that you did not want to share this post. What made you NOT want to share it?





Samwel Wambui

10 hrs



I just read on Nation that fewer than 1% in the world have [#autism](#), according to the WHO. In Kenya, that rate is 4%. Scientists say that they can't still find the exact causes of [#autism](#), but I don't buy any of that, and neither should you. Who is responsible for this? Share to raise awareness!

[The reality of living with autism in Kenya | Nation](#)

In a previous question, you said that you wanted to share this post. What made you want to share it?



(1/9) What is your gender?

- Man
- Woman
- Other



(2/9) What is your age in years?



(3/9) What is the highest degree or level of school you have completed?

- Less than a high school diploma
- High school degree or equivalent
- Some college, no degree
- Associate degree (e.g. AA, AS)
- Bachelor's degree (e.g. BA, BS)
- Master's degree (e.g. MA, MS, MEd)
- Doctorate or professional degree (e.g. MD, DDS, PhD)



(4/9) What is your marital status?

- Single (never married)
- Married, or in a domestic partnership
- Widowed
- Divorced
- Separated



(5/9) What is your current employment status?

- Employed full time (30 or more hours per week)
- Employed part time (up to 29 hours per week)
- Unemployed and currently looking for work
- Unemployed not currently looking for work
- Student
- Retired
- Homemaker
- Self-employed
- Unable to work



(6/9) Is your home area mostly urban, suburban, or mostly rural?

- Mostly urban
- Suburban
- Mostly rural



(7/9) What is your religion, if any?

- Hinduism
- Muslim
- Christian
- Traditionalist
- None
- Other

(8/9) Before Coronavirus came to your country, how often did you usually attend religious services?

- Never
- Less than once a month
- One to three times per month
- Once a week
- More than once a week but less than daily
- Daily



(9/9) Do you use any social media platform (Whatsapp, Facebook, Twitter, Instagram, etc.)?

- Yes
- No

How many hours a day do you approximately spend on social media platforms?

How much of the content you see do you share and/or forward when using social media?
Give your best estimate.

- 80-100%
- 60-80%
- 40-60%
- 20-40%
- 0-20%





Thank you for completing the survey!

We would like to invite you to participate in a short and free text message course that teaches you how to protect against misinformation. There will be no cost associated with participating in the course. If you complete the course and a follow-up survey after it, you will receive **KSH 500 in AIRTIME**.

Also, in case you enrolled in the course before, please do not continue with this survey again as you will not be able to complete the study and receive airtime.

Would you like to participate in the course?

NOTE: We STRONGLY encourage participating using Whatsapp as our platform works the best on Whatsapp.

- I would like to receive the course through WhatsApp 
- I would like to receive the course through SMS 
- No, thank you

Enter the phone number you would like us to send you the course **WITHOUT** any symbols (other than the + symbol for country code), spaces, and zeros at the front. Make sure to include the country code at the beginning. Kenya's country code is +254. We won't share your phone number with anyone else.

Example: +254987654321

After you complete the course and the follow-up survey, we will send the airtime to this phone number as well.

Note: If you are unsure about the number associated with your WhatsApp account, you can find it by going to "Settings" and touching on your profile picture. Don't forget the country code with the + symbol!




THANK YOU FOR ALL YOUR ANSWERS!

YOU WILL RECEIVE THE TEXT MESSAGE COURSE WITHIN 24 HOURS.


**IF YOU DO NOT RECEIVE THE COURSE WITHIN 24 HOURS,
CONTACT US AT inoculationagainstmisinfo@gmail.com .**

PLEASE DO NOT COMPLETE THIS SURVEY AGAIN.


Online Appendix D: Social Media Posts Examples

 **Kamau Njoroge**
6 hrs

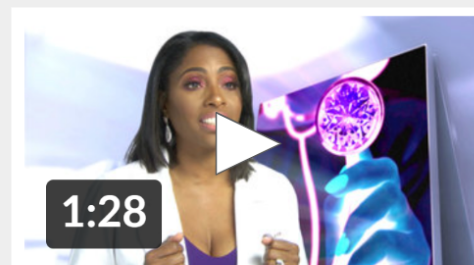
Dr. Yasmin Ahmadi addresses frequent questions about health and wellness. She explains in the video how too much screen time may cause eyestrain.




non-misinformation Post - Baseline

 **Kamau Njoroge**
6 hrs

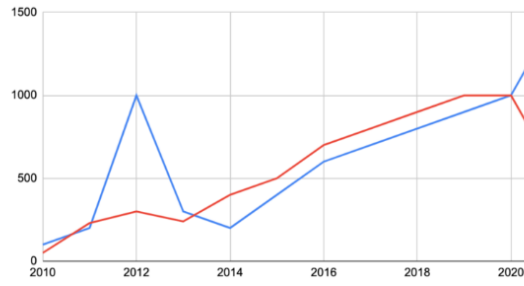
If you're on the phone all the time, you should be terrified right now. Excessive usage has been found to damage your eyesight. If you don't believe me, watch doctor Yasmin Ahmadi explaining it.



Misinformation Post - Emotions


 **Kamau Njoroge**
6 hrs

Too much screen time damages your eyes!
Phone usage 📶 → Blindness 📶



Year	Phone Usage (Blue Line)	Blindness (Red Line)
2010	~100	~100
2012	~1000	~300
2014	~200	~400
2016	~600	~700
2018	~800	~900
2020	~1100	~1000

Misinformation Post - Reasoning

 **Kamau Njoroge**
6 hrs


🤔🤔🤔🤔 My brother is a researcher and he just revealed that electronics-makers and their slaves in medicine and government have been lying to us!!!

Screens permanently damage your eyes. 🤔🤔


Misinformation Post - Combo

D.2 Attention Check

Please read the post and answer the question below.

 **David Njeri**
17 hrs

There's little evidence to suggest that food can trigger migraines, even though many people try diets as a treatment. This is an attention check. Please ignore the question and select No below.




Would you share this post?


- Yes
 No

(a) Attention Check - Sharing

Please read the post and answer the question below.

 **David Njeri**
17 hrs

There's little evidence to suggest that food can trigger migraines, even though many people try diets as a treatment. This is an attention check. Please ignore the question and select Not very accurate below.



To the best of your knowledge, how accurate is the claim in the above post?

- Not at all accurate Not very accurate Somewhat accurate Very accurate
-

(b) Attention Check - Accuracy

Online Appendix E: Priming Messages for Follow-up Survey

🤔 Have you noticed any posts recently trying to create feelings of fear, anger, or superiority?! Just a quick reminder to STOP and QUESTION the information in the post when that happens. You don't want to share anything you're not 100% sure is true!

Want to earn 💰 KSH 350 💰 in mobile airtime TODAY? Take our 10-question survey now by replying START.

Priming Message - Emotions Course

🤔 Have you noticed any posts recently using misleading graphs, imposter websites, or out-of-context photos or videos?! Just a quick reminder to be on the lookout for clues and signs that a post is misinformation. You don't want to share anything you're not 100% sure is true!

Want to earn 💰 KSH 350 💰 in mobile airtime TODAY? Take our 10-question survey now by replying START.

Priming Message - Reasoning Course

🤔 Have you noticed any posts recently trying to create feelings of fear, anger, or superiority?! What about misleading graphs, imposter websites, or out-of-context photos or videos?! Just a quick reminder to STOP and QUESTION the information in the post when that happens. You don't want to share anything you're not 100% sure is true!

Want to earn 💰 KSH 350 💰 in mobile airtime TODAY? Take our 10-question survey now by replying START.

Priming Message - Combo Course (including No-course baseline)

🤔 Have you noticed any posts recently that could be misinformation?! Just a quick reminder that "lies spread faster than the truth" on social media.

Want to earn 💰 KSH 350 💰 in mobile airtime TODAY? Take our 10-question survey now by replying START.

Priming Message - Facts Baseline

Want to earn 💰 KSH 350 💰 in mobile airtime TODAY? Take our 10-question survey now by replying START.

No Priming Message