

“I can say I’m John Travolta ... but I’m not John Travolta.”* Investigating the Impact of Changes to Social Media Verification Policies on User Perceptions of Verified Accounts

Carson Powers*, Nickolas Gravel*, Christopher Pellegrini*, Micah Sherr**,
Michelle M. Mazurek[†], and Daniel Votipka*

*Tufts University; [†]University of Maryland; **Georgetown University
{carson.powers,nickolas.gravel,christopher.pellegrini,daniel.votipka}@tufts.edu
mmazurek@cs.umd.edu; msherr@cs.georgetown.edu


Abstract

Until recently, almost all social media platforms verified the identities behind notable accounts. Prior work showed users understood this process. However, Twitter/X’s switch to an open, less rigorous verification process represented a significant policy shift. We conduct a U.S. Census-representative survey to investigate how this and subsequent verification changes across social media impact users’ verification perceptions. We find most users generally recognize the changes to Twitter/X’s policy, though many still believe Twitter/X verifies account holders’ true identities. However, users are less aware of subsequent Facebook verification changes. We also find platforms’ verification differences do not impact user perceptions of posted content credibility.

Finally, we investigate various hypothetical verification policies. We find participants are more likely to perceive posts from verified accounts as credible when only notable accounts are eligible and government document review is required. Payment did not affect credibility decisions, but participants felt strongly that payment for verification was unacceptable.

*The full quote by a participant asked what verification policy changes they would suggest was, “I would require a photo ID. I can say I’m John Travolta and I can give you my email address (which can be almost anything) to confirm me, but I’m not John Travolta.”

1 Introduction

Most social media sites, such as Twitter/X¹, Facebook, TikTok, and LinkedIn, support some form of account verification. Each platform reviews accounts [41], then adds a badge (e.g., ) next to the *verified account’s* (VA)² username to signal the verification process has been completed. VAs were introduced to help users differentiate between accounts belonging to the entity named (often a celebrity or account of public interest) and parodies or impostors [71]. Twitter introduced VAs in 2009 following a rise in impostor accounts [64], and other platforms followed suit [22, 33, 41, 59, 61, 69]. With the rise of disinformation on social media, the value of determining a post’s true source is growing [28, 35–37, 48, 56]. This challenge is exacerbated in emergencies, when users look to social media for real-time information [6, 31, 40, 46, 73]. During terrorist and active-shooter events [3, 6] and natural disasters [40, 45, 52], users look to local authorities, such as police and fire departments, for safety information. Without rigorous account verification, users may trust false information with life threatening consequences [27].

While there is some evidence suggesting users equate account verification with credibility [43], other work has shown, in isolation, users correctly understood the verification badge only indicated authenticity [72]. However, recent changes to verification policies may muddle verification’s purpose. First, the social media ecosystem has splintered, with new and niche platforms growing (e.g., TikTok, Truth Social, etc.). While verification is similar across platforms, some subtle differences should impact the correct interpretation of VAs.

Additionally, some of the largest existing platforms have made significant policy changes. Most notably, Twitter dramatically changed its verification policy after being acquired by Elon Musk in October 2022. Prior to the purchase, Twitter verified notable users’ accounts (e.g., celebrities and public

¹Since Twitter’s rebranding to X occurred after our survey, we will use “Twitter” in the remainder of the paper.

²Terms differ by platform. For consistency, we refer to accounts that have undergone some form of authentication as verified accounts (VAs).

figures or organizations) by requiring proof of identity via a government-issued ID [71]. Twitter then made verification available to any user for a monthly \$8³ subscription fee, and swapped government ID for a verified phone number [70]. This transition was tumultuous, with abrupt changes regularly covered in the media [9, 20, 29, 32, 44, 47]. Some users took advantage of the new policy to establish impostor accounts [49, 63]. To less fanfare, Facebook also adjusted its verification policy, allowing anyone to obtain a VA for a fee, but maintaining the requirement for ID verification, and LinkedIn made verification slightly more open without adding a fee.

We seek to assess the impact of these policy changes on user perceptions of VAs, as well as how users think verification policies *should* work. Towards that goal, this paper considers three research questions:

RQ1: What are the verification policies used by popular social media platforms and how have they changed over time?

RQ2: What do users think account verification entails? How does it impact perceptions of posted content credibility?

RQ3: How would potential changes to verification policies impact user perceptions of posts from verified accounts and user perceptions of the policies?

RQ1 seeks to understand the VA ecosystem. Due to the fractured landscape, perceptions may vary depending on the platforms used. With the volume of media coverage and rapid policy-making during the Twitter transition, user perceptions may represent a snapshot in time, rather than an accurate depiction of current policy. To understand the impact of these changes, we must first enumerate verification policies.

To address RQ1, we captured the verification policies of eight popular social media sites from April 2022 to August 2023, noting any changes. After enumerating verification policies, we conducted a controlled experiment—using a vignette-based survey of 1600 U.S. Prolific users—to address RQ2 and RQ3. Participants were first shown two mock posts containing contradictory information and asked to indicate which they perceived as more credible, to test the VA’s impact on their assessment of relative credibility when presented with information from similar accounts—a common challenge when assessing information during emergency events. We varied the platform (Twitter vs. Facebook) and asked participants to indicate how they believed their assigned platform defined verification. Then, we presented participants with a new verification policy and asked them to reevaluate the previously shown mock posts with this new policy in mind. We also asked participants their perceptions of the new policy.

Participants’ understanding of Facebook’s and Twitter’s verification policies was mixed, and they were more likely to correctly perceive Facebook’s policy as requiring identity verification. Participants correctly indicated Twitter’s policy was

open to anyone for a fee. This seems to indicate users have better understood the Twitter policy over time, compared to a similar survey conducted earlier by Xiao et al., which asked participants to identify features of verification [78]. However, participants seemed unaware of Facebook’s policy, with many still believing verification was free and only for notable accounts. This is likely due to the newness of Facebook’s policy change and lack of broad media coverage.

We did not observe differences in participants’ assessments of posted content credibility between assigned platforms. However, after providing participants with a verification policy, they were more likely to find posts from the VA credible when government ID was required and only notable accounts were verified. Participants also perceived these policies as more acceptable (matching Xiao et al. [78]). This difference between initial assessment and re-assessment after reviewing a verification policy suggests participants do not consider the details of the policy fully when assessing posts from VAs.

Finally, while participants strongly disliked paying for verification—corroborating Xiao et al. [78]—payment did not impact participants’ credibility decisions before or after reviewing the verification policy. While this indicates verification payment has no direct impact on user assessments of credibility of VAs’ posts, the strong dislike of the policy may have downstream impacts that should be considered in future work, especially as several participants reported no longer trusting any verification provided by Twitter.

2 Related Work

Credibility of Online Content. A large body of work has investigated factors affecting user perception of online content credibility. Wineburg et al. surveyed students to assess ability to judge online sources’ credibility [76]. Fogg et al. found the “design look” of a website impacts perceived credibility [18]. Hilligoss and Rieh found users are more likely to find information legitimate when the source appears “official” [26]. Hassoun et al. performed a qualitative analysis of Gen Z’s evaluation of online information, finding three “trust heuristics”: credible information was easily accessible, neutral in tone, and “felt right.” Their participants reported using number of likes and comments as a form of “crowdsourcing credibility” [24]. This mirrors previous findings that users are more likely to perceive information as credible when they believe others perceive it as credible [8, 17, 21, 26, 66], an effect called the *endorsement heuristic*. Familiarity with a source also increases perceived credibility, known as the *reputation heuristic* [42]. We build on prior work, focusing specifically on social media platforms and the effect of verified indicators.

Verification’s Impact on Social Media Post Credibility. The verified indicator’s purpose is to affirm an account holder’s identity, not signal posted content credibility. However, humans’ reliance on trust heuristics may lead to an indirect

³\$12 if signing up in-app to account for Apple’s/Google’s service charges.

effect on perceived credibility, which may explain conflicting evidence whether users separate *authenticity* and *credibility*.

Early work by Morris et al. suggested the verified indicator has a high impact on users’ evaluation of credibility [43]. However, their work asked participants to list features they consider when deciding if a tweet is credible, which measures the *conscious* impact of verification badges, not the *behavioral* impact. Conversely, Vaidya et al. conducted a large-scale controlled experiment, measuring the verified indicator’s effect on participants’ perceptions of post credibility. They found users understood verification indicated the account holder was who they said they were, but does not add credibility to the post [72]. Similarly, in 25 interviews with social media users about fake news, Geeng et al. found most users do not conflate authenticity with credibility [21]. Dumas and Stough conducted a consumer-behavior study where participants were shown influencer-posted content. They found consumers associate VAs with celebrity more than source credibility [12]. In this paper, we seek to assess whether user perceptions have changed due to changes to social media verification policies and expand beyond Twitter to consider other platforms.

Most similar to our work, Xiao et al. investigated user understanding of verified indicators on Twitter, Facebook, and TikTok in the wake of Twitter Blue [78]. They surveyed social media platforms and identified the dimensions of each verification policy. Using these, they surveyed 299 U.S. adults asking their definitions for verification and whether they found Twitter’s policy acceptable. They found participants were more likely to indicate payment was required for Twitter, as opposed to other platforms, but most continued to *incorrectly* assume Twitter verified the identity of users with verified indicators. They also observed users disliked Twitter’s policy because it does not verify user identity and requires payment. We build on this study in several ways. First, we conducted a more in-depth review of social media platforms by investigating Musk’s Twitter posts, which provide valuable context, and monitoring policies over a longer period, which allowed us to capture policy changes by Meta and LinkedIn. Next, our survey captured a snapshot in time after Meta’s policy changes, allowing a useful comparison over time between the works. We also go beyond understanding how users define verification by measuring how verified indicators impact perceptions of post credibility. Finally, we conduct between-subjects comparisons, randomly assigning participants to define verification for specific platforms, instead of asking for general definitions, and test several possible policy designs for their impact on participant post credibility decisions and policy acceptability. This gives us a more nuanced picture of the changing landscape of VAs and its impact on user behaviors.

3 Verification Policy Review

To address RQ1, we reviewed verification policy changes across eight popular social media platforms from April 2022

to August 2023. We outline our collection and review process and describe changing landscape of social media verification.

3.1 Data Collection and Analysis

We collected verification policies from seven of the top eight social media platforms Americans reported getting their news from in 2022 [7], i.e., Twitter, Facebook, TikTok, Snapchat, LinkedIn, Instagram, and YouTube. We excluded Reddit, which does not support account verification, but included Truth Social to represent small, niche platforms.

For each platform, we captured the verification policy on April 14, 2022 and all subsequent policy changes until August 25, 2023. April 14 marked Elon Musk’s expression of interest in acquiring Twitter. This date served as a significant marker for our analysis, as it potentially influenced changes in the verification policy landscape. We continued monitoring platform policies until our final participant completed our survey (see Section 4) to ensure we captured changes that could affect user perceptions. For brevity, details about our web scraping process are included in Appendix B.

We also manually reviewed all of Musk’s personal tweets about Twitter’s verification policy during this period. Musk regularly made policy pronouncements publicly, which drove news coverage [32, 62] and may have influenced perceptions.

To identify common themes across verification policies, we performed an inductive thematic analysis, allowing policy dimensions to arise from the data [65]. Two researchers collaboratively reviewed the initial policies for each platform and subsequent changes as they were collected. Codes were then discussed with the full research team until full agreement was reached. Because we only sought to identify themes and do not attempt to use results for quantitative comparison, we did not assess inter-rater reliability [38].

3.2 Results

We observe several independent *dimensions* of social media verification policy: who can be verified (Eligibility), how accounts are verified (Verification Method), whether users pay a fee, requirements to prevent “deception,” and required activity history. Table 1 summarizes the reviewed policies, including any changes occurring during our review.

Further, we observe three distinct *time periods* of social media verification policy:

Before Musk’s Twitter takeover (Period 1). From the start of our review (April 14, 2022) until Musk’s takeover of Twitter (October 27, 2022), the policies of all eight social media platforms were similar. All allowed verification only for “Notable” users (e.g., celebrities, journalists, public figures). They required users provide government documents to prove identity and did not charge for verification. There was some variation in what platforms considered “deceptive.” These policies prevent accounts from changing their account information

Platform	Icon	Eligibility	Ver. Meth.	Payment	Non-Deceptive	Active
Twitter [70, 71]	✓	Notable → Open	Gov ID → Phone	Free → Paid	No profile changes, ¹ spam, misleading behaviors, or platform manipulation	Active past 30 days
Facebook [41]	✓	Notable → Open	Gov ID	Free → Paid	No profile changes, ¹ unique	Prior posting history
Instagram [41]	✓	Notable → Open	Gov ID	Free → Paid	No profile changes, ¹ unique	Prior posting history
TikTok [69]	✓	Notable	Gov ID	Free	No profile changes ¹	Logged in past 6 months
Snapchat [59]	★	Notable	Gov ID	Free	No misleading behaviors	Regularly post content
LinkedIn [33]	✓	Notable → Open	Gov ID ²	Free	No profile changes	-
YouTube [22]	✓	Notable	Gov ID ³	Free	No profile changes ¹	Regularly post content
Truth Social [61]	✗	Notable	Gov ID	Free	No misleading behaviors	Regularly post content

¹ All platforms restricted VAs from changing their username. Some also prevented changes to other profile data, such as profile photos and bios.

² LinkedIn’s verification is only available to US users (through the CLEAR ID program) or employees of companies participating in LinkedIn’s company email verification or Microsoft’s Entra Verified ID programs.

³ YouTube does not verify documentation by default, but reserves the right to request additional documentation if necessary.

Table 1: Summary of verification policy dimensions and verified indicators per platform. → indicates a change in the policy during our review with the left hand side indicating the policy at the start of our review and the right hand side showing the final policy.

(e.g., username), having usernames similar to other accounts, posting spam, or attempting to manipulate the platform.

Musk acquired Twitter (October 27, 2022; Period 2). Musk made sweeping verification policy changes by introducing Twitter Blue on November 9, 2022. This program opened verification to any user, removed user identity checks, and required payment [70]. Musk argued open verification would improve conversation quality [13] and reduce bots by creating a barrier to entry [14, 15]. These changes faced broad criticism [9, 20, 29, 32, 44, 47], and verified impostor accounts quickly appeared [49, 63], indicating the changes did not produce Musk’s desired effect [67].

Twitter paused Twitter Blue on November 11, 2022 and reintroduced it on December 12, 2022 with modified eligibility requirements to limit impostors. Specifically, users were required to verify a working phone number and must have been active 30 days before verification.⁴ Twitter also introduced government (✓) and company (✓) badges which were only available to organizations fitting these descriptions.

Potentially adding to user confusion, users verified under Twitter’s original verification policy (Twitter Legacy) maintained their verified indicator. Verification of Twitter Legacy and Blue accounts was indistinguishable when looking at individual posts. The only distinction was an indicator on the Twitter Legacy accounts’ profile pages. The Twitter Legacy policy remained in effect until April 1, 2023 [47].

While not directly related to the verification policy, Twitter also began prioritizing posts by VAs (January 5, 2023) [74]. Twitter argued this was to ensure users are most likely to see “content that is relevant, credible, and safe,” implying a link between verification and credibility.

During this period, all other platform policies were stable.

⁴The policy initially added a 90-day activity period on November 24, 2022, but this was relaxed to 30 days prior to Twitter Blue’s restart.

Meta and LinkedIn alter policies (February 20, 2023; Period 3). Meta, the parent company of Facebook and Instagram, announced Meta Verified [41]. Like Twitter Blue, this subscription-based verification program was open to all users and required payment. However, Meta continued to require government ID for verification—the most significant difference between Twitter’s and Facebook’s final policies.

On April 12, 2023, LinkedIn also opened verification eligibility beyond notable users [33]. LinkedIn began allowing U.S. users to verify their identities through the CLEAR ID program and verified users with certain corporate email addresses or through the Microsoft Entra Verified company ID program. While not available to all users, it is more open than previously, and follows Meta’s example of maintaining identity verification while increasing eligibility.

Our identified dimensions of verification align with those outlined by Xiao et al.’s prior review [78], though our results capture changes to Meta’s and LinkedIn’s policies that occurred after their review. Our full dataset of policy changes can be found at https://osf.io/a9y3j/?view_only=d2608dffe87f40c09885c4e55637ddeb.

4 Survey Methods

Using the policy dimensions identified in Section 3, we developed an online survey to test participants’ understanding of platform policies (RQ2) and their preferences for each policy dimension (RQ3).

4.1 Survey Design

Figure 1 shows the stages of our online survey, which we describe below in turn.

Consent (Part A). We began with a consent form describing the study, potential risks, and data protection procedures. To




Figure 1: Sections and flow of the user study.

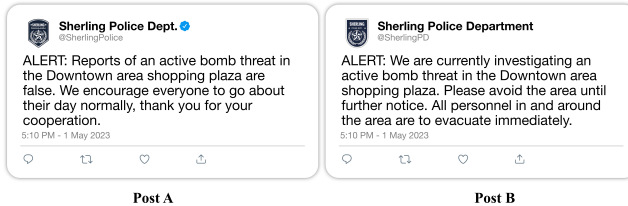


Figure 2: Example Police/Declarative/Twitter condition posts.

avoid priming for the verified indicator, which users might otherwise ignore in practice, we used deception when describing the study’s purpose, indicating it was to understand how users assess social media posted content credibility.

Implicit effect of verified indicator (Part B, RQ2). Next, participants were shown a pair of posts reporting contradictory information, both from accounts presenting as authorities on the subject. Figure 2 shows an example pair of posts. Posted content details, such as whether they included a verified indicator and the platform for which they were formatted (Twitter or Facebook) varied per condition (see Section 4.2). Participants were asked to indicate which posted content was more likely correct, on a five-point Likert scale. Because the contradictory posts cannot both be true, participants must make some assessment (potentially based on the verified indicator) about account identity to determine which is more credible.

Explicit effect of verified indicator (Part C, RQ2). Next, we asked participants whether the verified indicator affected their posted content credibility choice, on a four-point Likert-type scale from “No effect” to “Major effect.” To compare the verified indicator’s effect to other account features, participants were asked the same question about the account’s picture, name, and handle.⁵ The order of account feature questions was randomized to avoid ordering effects [55].

Participants’ verification definitions (Part D, RQ2). We then asked participants to define verification to investigate how they understand verification and if this varies by platform.

Assigned verification policy perceptions (Part E, RQ3). We gave a mock verification policy and asked participants to assume their condition-assigned platform adopted this policy. We asked whether they believed it was “acceptable for verifying account owner identity” on a 5-point Likert-type scale

⁵The account handle question was only included for participants in the Twitter condition because Facebook accounts do not have this feature.

from “Unacceptable” to “Acceptable.” We also asked them to provide one modification (i.e., addition, deletion, change) to improve the policy. This open-ended question was intended to capture the policy elements participants prefer and prioritize, including those not used on social media platforms.

Credibility perceptions after policy priming (Part F, RQ3).


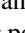
In Part F, we showed participants the original contradictory posts together with their assigned mock verification policy. Then, we repeated Part B’s question, asking participants to choose which posted content was more credible, this time assuming verification via the given mock policy. Next, we asked participants to assume a friend was unsure which posted content was more credible, and tell us what advice they would give to help the friend decide. This open-ended question captured an additional perspective into participants’ credibility assessment. This section included an attention check to identify and remove inattentive respondents [39].

Social media use (Part G), Security attitudes (Part H), and Demographics (Part I). We concluded with questions about our participants’ background and demographics. We asked about their social media use for the two platforms tested, as well as more generally. Participants completed Faklaris et al.’s SA-6 scale [16] to assess their computer security practices.

Debrief (Part J). Because we used deception, we debriefed participants about the study’s true nature, providing Twitter’s and Facebook’s verification policies and links to best-practice guidelines for assessing posted content credibility [34, 57, 58].

4.2 Conditions

Each participant saw two contradictory posts (Parts B and F) and a mock verification policy (Parts E and F). We describe the possible posts and policies defining each *condition*.

Posted content variables. To test the verified indicator’s effect, we created four posted content pairs. First, we varied the platform. One of our research questions (RQ2) is whether users perceive differences in verification policy between platforms and how this impacts VA credibility perceptions. For this dimension, participants were shown posts using Twitter or Facebook visual cues. This included the posted content design, verified indicator shown (i.e.,  vs. ) and terminology in survey questions (e.g., “Please answer the following questions considering the two *Twitter* posts above⁶”). We chose these platforms because Twitter changed its verification policy most significantly (see Section 3) and Facebook was the most popular platform with a comparable modality (i.e., YouTube, Instagram, TikTok are mostly image and video-based).

Second, we varied the posted content. Prior work showed content affects users’ credibility perceptions [72], so we test multiple content types to avoid bias from a single type.

⁶Emphasis not included in survey.

One pair of posts describes an alleged bomb threat (*Police*), as posted by different accounts (Sherling Police Dept. @*SherlingPolice* or Sherling Police Department @*SherlingPD*) claiming to be the same entity. One post claims the threat is false; the other asserts it is true. The second pair (*Coffee*) appear to be posted by medical doctors (Dr. Samuel Smith, M.D. @*DrSmithMD* or Dr. Alexander Kim, M.D. @*DrKimMD*). The posts contradict about a link between coffee consumption and risk of a disease. Table 2 details the posts. Combining platform and content options produced four posted content conditions. Participants were randomly assigned to one.

The named police departments, doctors, and diseases were fictional to eliminate prior knowledge bias. We avoided hot-button or political topics to prevent polarization effects [30]. We chose topics of general importance, where people must rely on expert insights. We chose to use authoritative accounts, as accounts like these could be verified or unverified under all policies reviewed in Section 3, creating a range of reasonable justifications participants could come to in their decision-making. Prior work showed users are more likely to find authoritative accounts credible [72], so we only used authoritative accounts to control for this effect.

To control for potential bias toward declarative or contradictory statements, we randomized which account was verified. We randomized the order the declarative and contradictory posts were shown, to control for ordering effects [55]. To control for other possible credibility indicators, other posted content elements (author profile image, retweets and likes counts, and time since publication) were held constant. Previous research showed these elements significantly affect user perception of posted content credibility [43].

Policy variables. After asking about the pair of posts, we presented participants with mock verification policies to observe how varying policy definitions affect their perception of the verified indicator (RQ3). The policies had three variables, representing the three dimensions we observed multiple platforms change in our policy review (Section 3). Table 2b gives the policy text shown for each condition. First, we varied who can be verified (Eligibility). The policy was either *Open*, meaning anyone can apply, or *Notable*, meaning only well-known individuals and organizations are eligible. Next, policies varied in how accounts are verified (Verification Method). That is, accounts must either confirm an email or phone number (*Phone*) or provide government-issued ID (*Gov ID*). Finally, the policy specified whether verification required *Payment*. We used a full-factorial variable combination to create eight policies. Participants were randomly assigned a policy independent of their post and platform condition.

4.3 Recruitment

We conducted our survey on Prolific, a research recruitment service providing high-quality samples [50, 68]. We limited

Content	Position	Posted content Text Summary
Police	Decl.	ALERT: We are currently investigating an active bomb threat in the Downtown area shopping plaza. Please avoid the area. . .
	Cont.	ALERT: Reports of an active bomb threat in the Downtown area shopping plaza are false. . .
Coffee	Decl.	Individuals who consume more than three cups of coffee per day may have a higher risk of developing end thrombocytosis.
	Cont.	There have been no research studies that have established a link between coffee consumption and end thrombocytosis.

(a) Posted Content Variables

Dimension	Option	Policy Text
Eligibility	Open	Any user on the platform is allowed to apply for verification
	Notable	Only well known, high-profile individuals and organizations are allowed to apply for verification
Verification Method	Phone	Accounts are required to confirm a phone number or email with the platform
	Gov ID	Accounts must submit government-issued identification matching the name of the account
Payment	Paid	Accounts pay a monthly subscription fee to maintain their verification checkmark
	Free	Accounts do not pay any fee to maintain its verification checkmark

(b) Policy Variables

Table 2: Summary of (a) posted content and (b) policy conditions. There were four posted content and eight policy conditions, resulting in 32 total conditions after a full-factorial combination.

participation to Prolific users at least 18 years old and located in the United States. We used Prolific’s census-representative sample feature [53] to ensure a U.S. population-representative distribution by age, gender, and ethnicity. Survey completion time averaged 8.2 minutes, and we paid participants \$2.

4.4 Pilot

We piloted the survey with nine participants—drawn from a convenience sample, selected for varying social media familiarity. Pilot participants were asked to “think aloud” while answering questions. We iteratively updated the survey for clarity after each pilot until further changes were unnecessary.

We also tested a third content type about an E.Coli outbreak in lettuce. We recruited 50 participants on Prolific and assigned them randomly to one of the three content types to test whether any content type behaved unexpectedly (e.g., prior experience bias or unexpected relationship with current events). We did not observe unexpected responses, but saw similar results between the E.Coli and Coffee conditions.

Therefore, we dropped the E.Coli condition to increase our analysis power by recruiting more participants per condition.

4.5 Data Analysis

Quantitative analysis. To test verification’s effect on participants’ posted content credibility perceptions before and after stating a policy, and to assess participants’ perception of policy acceptability, we used ordinal logistical regressions.

For the two posted content credibility perceptions questions, the outcome variable is a 5-point Likert-scale response regarding which post was correct (Part B and Part F, respectively). Each response was modified to indicate whether the participant perceived the account with or without the verified indicator as correct, to allow for comparisons; e.g., if a participant shown the posts in Figure 2 selected “Definitely A” from the possible options, because A was the VA, their response was modified to “Definitely the VA.” For the policy acceptability regression, the outcome variable was the participant’s response to the policy acceptability question in Part E.

In each regression, we include the assigned condition’s three elements (platform, content, and position) as explanatory variables. For the policy-related regressions (Part E and Part F), we added the policy variables (Eligibility, Verification Method, and Payment). In all regressions, we include demographic explanatory variables (age, gender, education), amount of time spent using Twitter and Facebook, number of social media platforms used, and SA-6 scores. Table 6 in Appendix D summarizes the variables included per regression.

To select a parsimonious model without overfitting, we constructed initial regression models using all possible explanatory variable combinations. We selected models with the minimum Bayesian Information Criterion, appropriate for testing goodness-of-fit [54, 60].

We also examined the explicit impact of verified indicator on credibility perceptions. We compared responses regarding the verified indicator’s impact between Twitter- and Facebook-assigned participants using a Pearson’s χ^2 test, appropriate for categorical data [19]. Next, we compared responses across the four⁷ account features (verified indicator, account username, photo, and handle) using non-parametric, repeated measures tests, appropriate for multiple Likert-scale responses per participant. We began with an omnibus Friedman test across features to control for Type I error; if the result was significant, we applied the Wilcoxon signed-rank test to planned pairwise comparisons of the verified indicator with every other feature [75]. Comparisons were across content conditions.

Qualitative analysis. We used iterative open coding to analyze free-response questions [65]. As our questions were all related to VAs and verification policies, similar to the free-response questions in Vaidya et al. [72], we began with their

⁷Three for participants assigned Facebook because they were not shown a user handle.

codebook. However, as verification policies have changed, we allowed additional codes to arise inductively. Three researchers extended the initial codebook collaboratively by reviewing 10 responses. Two researchers independently coded additional responses in rounds of 100, updating the codebook incrementally. After rounds, the coders met, assessed inter-rater reliability using Krippendorff’s alpha [25], and resolved coding differences. After two rounds (200 responses), the coders achieved $\alpha = 0.80$, which represents acceptable agreement. The remaining 1386 responses were divided evenly and coded separately by the two coders [25]. Finally, the two researchers performed an axial coding to identify relationships between codes and produce higher-level groups [10, pg. 123-142]. Appendix F gives the final codebook.

To compare initial verification definitions between participants shown Twitter and Facebook posts, we perform Pearson’s χ^2 tests, appropriate for categorical data [19]. For each higher-level code group, we compare a code’s presence from this group between Twitter- and Facebook-assigned participants. Because this requires multiple testing, we apply a Benjamini-Hochberg correction to adjust p -values [5].

4.6 Ethical Considerations

Our institution’s IRB approved this study. We obtained informed consent prior to the survey. Because we used deception in our study description, we concluded with a debrief and asked participants to re-consent. To avoid response coercion, participants were told they would be paid for completing the survey even if they refused consent, but their response would be deleted. Three participants withdrew after the debriefing.

Responses through Prolific are provided pseudonymously, with only the participant’s Prolific ID identifying their response. We did not request additional identifying information.

4.7 Limitations

We presented mock posts, as this provides the control needed to reason about specific variables’ effects on credibility perceptions. However, we are unable to capture other credibility perception influences, such as the posting author’s reputation, the participants’ relationship with the author, or the participants’ relationships with others who interact with the posted content (e.g., liked or shared). The types of content and other metadata we test are also limited, meaning we are unable to comprehensively test these factors’ influence on posted content credibility. This is an inherent tradeoff to limit the study’s scope to a reasonable condition set. We believe our conditions are sufficient to target our study’s research questions.

The study’s setting also differs from the real world. Participants may have spent more time reviewing our contradictory posts than when casually browsing social media feeds. Also, presenting contradictory information side-by-side is not representative, as these posts would be interspersed with other

posts. Our results are indicative of a best-case situation where users carefully consider all relevant information, which is likely closer to the truth in emergency situations when finding good information is safety-critical and social media is saturated with posts about an ongoing event.

For open-response questions, we give the percentage of participants who stated each theme. However, not mentioning a theme does not indicate disagreement. Participants may have failed to state an idea or considered other thoughts more relevant. Our open-response results should be viewed as a measure of what was “front of mind” when answering.

Even though we used Prolific’s census-representative sample feature, Prolific users are often more knowledgeable regarding privacy and security and more likely to use multiple social media platforms [68], which may impact generalizability. To account for these differences, we controlled for social media use and security attitudes in our regressions.

As these limitations apply across all conditions, we focus primarily on between-condition comparisons.

5 Survey Results

The majority of our key findings are taken from our regression analyses over initial perceived correctness (Table 4a), perceived correctness after proposing a new policy (Table 4b), and perceived policy acceptability (Table 5). Only variables in the final selected model are shown (as groups of rows). We give the base case first for categorical variables. We selected base cases expected to correlate with the lowest levels of VA perceived correctness and policy acceptability.

For categorical variables, OR is the odds ratio of the outcome (e.g., acceptability) increasing one Likert-scale unit when switching from the base case to the given parameter level. For numeric variables (e.g., SA-6), OR is the odds the outcome increases one Likert-scale unit for each one-point increase in the numeric variable. For example, the OR for *Police* in Table 4a indicates a participant assigned *Police* instead of *Coffee*—holding all other variables equal—would be $1.57 \times$ as likely to increase one unit in perception that the VA posted the correct message. Because this effect is greater than one, participants are more likely to report the VA as correct for *Police* than *Coffee*. *Police*’s confidence interval (CI) indicates that if we ran the study many times, we would expect 95% of runs to produce ORs between 1.31 and 1.87. The p -value (< 0.001) is less than our significance threshold ($\alpha = 0.05$), indicating a significant difference between *Police* and *Coffee*.

5.1 Participants

1739 participants attempted and 1660 completed the survey. We removed 27 who failed the attention check, 30 who gave nonsensical or obviously AI-generated responses to open-ended questions, and 3 who withdrew after the debrief. Our final dataset contains 1600 responses (50+ per condition).

Metric	%	Metric	%
Age		Education	
18-29 years	23.8%	H.S. or below	13.0%
30-49 years	34.9%	Some college/ Assoc.	32.9%
50-64 years	28.9%	B.S. or above	53.9%
65+ years	12.4%	Prefer not to respond	0.3%
Platform w/Account		Social Media Use	
Facebook	82.2%	<30 mins daily	19.1%
YouTube	78.9%	30 mins-1 hr daily	30.6%
Instagram	68.7%	1-2 hrs daily	28.7%
Twitter	66.7%	2-4 hrs daily	16.6%
LinkedIn	42.7%	5-6 hrs daily	3.3%
TikTok	37.0%	>6 hrs daily	1.6%

Table 3: Participant demographics. Percentages may not add to 100% due to non-response or selection of multiple options.

Table 3 summarizes participant demographics. Additional demographics are reported in Appendix D. Our participants’ gender and income were similar to the 2020 U.S. Census [1]. Participant ethnicities were similar to the U.S. Census, though White participants were overrepresented and Latino/a participants were underrepresented. Participants were more educated and younger on average than the U.S. population, though similar to estimated Twitter user demographics [77]. Participants’ average SA-6 score was 3.61, close to the average score from a U.S. Census-representative sample [16].

Participants most often had accounts with Facebook (82.2%), YouTube (78.9%), Instagram (68.7%), and Twitter (66.7%)—similar to other social media use surveys [4]. They most often used Twitter at least every other day (38.8%), with the majority using it at least once per week (64.9%), and many having no account (35.1%). Participants were more active on Facebook, with most using it at least every other day (56.1%) and only 19.8% not having an account. Facebook use did not vary significantly between participants assigned to the Twitter and Facebook conditions ($\chi^2 = 2.9, p = 0.566$). Twitter usage did vary between platform conditions ($\chi^2 = 9.6, p = 0.047$), but the effect size indicates little if any association ($\phi = 0.08$) [11, pg. 282].

5.2 Initial Impact of Verified Account (RQ2)

Here, we discuss participant perceptions of the contradictory posts’ credibility (Part B) and how they perceived the verified indicator impacting their decision-making (Part C) *prior* to being given a verification policy. Figure 3 summarizes initial credibility perceptions divided by experimental condition, and Figure 6 in Appendix D summarizes participants’ perceptions of the account features’ decision-making impact.

No difference between platforms. Across conditions, participant perceptions of the more likely credible post were evenly distributed. Participants most often indicated the VA was “Definitely” or “Probably” credible (43.9%). However, 32.1% indicated “Either the verified or not VA” was credi-

Variable	Value	Odds Ratio	CI	p-value
Content	Coffee	–	–	–
	Police	1.56	[1.31, 1.87]	<0.001*
Position	Contradict.	–	–	–
	Declar.	1.42	[1.19, 1.69]	<0.001*
Age	–	–	–	–
	+1	0.99	[0.98, 0.99]	<0.001*

– Base case (OR=1, by definition)

*Significant effect

(a) Initial Perceived Verified Account Correctness

Variable	Value	Odds Ratio	CI	p-value
Content	Coffee	–	–	–
	Police	4.13	[3.39, 5.02]	<0.001*
Availability	Open	–	–	–
	Notable	1.80	[1.50, 2.17]	<0.001*
Verification Method	Phone	–	–	–
	Gov ID	1.30	[1.08, 1.56]	0.005*
Facebook User	False	–	–	–
	True	1.54	[1.21, 1.95]	<0.001*
SA-6	–	–	–	–
	+1	1.21	[1.08, 1.36]	<0.001*

– Base case (OR=1, by definition)

*Significant effect

(b) Verified Account Correctness After Policy Given

Table 4: Summary of regression over participants’ VA correctness perception (a) before and (b) after being shown a specific policy. Pseudo R^2 measures for (a) were 0.01 (McFadden) and 0.04 (Nagelkerke), and for (b) were 0.07 (McFadden) and 0.17 (Nagelkerke).

ble and 24.1% chose “Definitely” or “Probably” the non-VA. Results were similar whether participants were assigned Twitter (43.8% VA, 33.3% either, 22.9% non-VA), or Facebook (44.0% VA, 30.8% either, 25.3% non-VA). The selected regression (Table 4a) did not include platform, indicating no observed statistically significant difference between platforms.

When asking participants directly about the verified indicator’s impact on their decision-making, responses again were split. A slight majority indicted it had no impact (52.0%), while 48.0% reported at least a “Minor effect.” Participants were statistically significantly more likely to rank the verified indicator’s effect higher than the account picture ($Z = 14.46, p < 0.001$) and handle ($Z = 7.31, p < 0.001$) according to Wilcoxon-Pratt signed rank tests. We did not observe a statistically significant difference between the verified indicator’s and account name’s perceived impact ($Z = 1.71, p = 0.087$).

Comparing platforms (Figure 6) there is no clear difference: 46.6% of Facebook-assigned participants reported at least a “Minor effect” versus 49.4% for Twitter. No statistically significant difference was observed ($\chi^2 = 4.82, p = 0.186$).

Content had the biggest effect. Participants shown the Police content were statistically significantly more likely to perceive the VA as credible ($OR = 1.56, p < 0.001$). If the VA posted

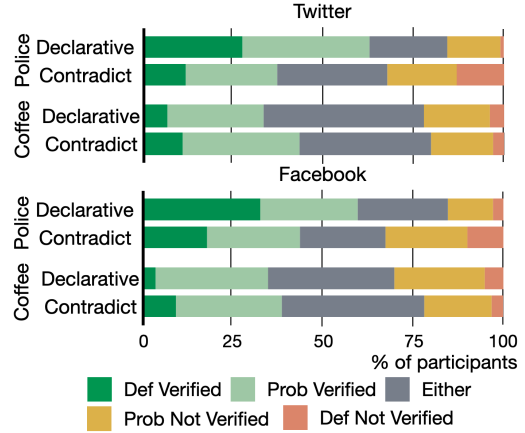


Figure 3: Likert-scale response showing whether participants perceived the VA as more likely credible, organized by assigned social media platform, content type, and the position taken by the VA.

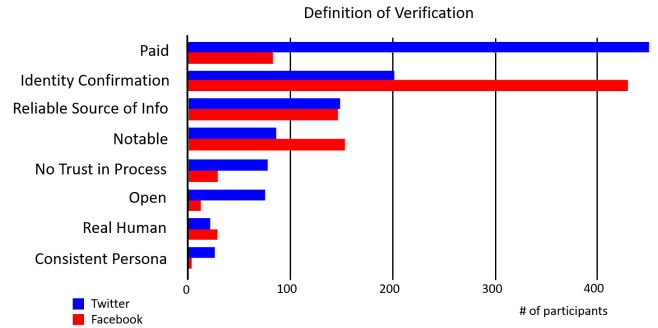


Figure 4: Participants’ verification definitions by platform.

the declarative statement (e.g., there was a bomb), participants were statistically significantly more likely to perceive the VA as credible ($OR = 1.42, p < 0.001$). This follows prior work [72], which showed content drives message credibility.

Age has some effect. Grouping participants by decade, we observed a downward trajectory in percentage of participants perceiving the VA as “Definitely” or “Probably” credible (52.1% of those under 30 to 24.1% of those over 70). Older participants were more likely to indicate “Either the verified or not VA” was credible (25.8% of the under 30s to 42.2% of the over 70s)—the correct response, as VAs do not necessarily post credible content. With each additional year, participants were $0.99\times$ as likely to find the VA more credible by one point ($p < 0.001$). When comparing an individual one standard deviation older (~ 15.75 years), we would expect them to be $0.85\times$ as likely to increase one point on the Likert scale. This contradicts prior results by Xiao et al. [78], who found no statistically significant relationship.

5.3 Verification Policy Definitions (RQ2)

Here, we discuss participants’ free-response verification definitions (Part D) prior to priming about a particular policy. These definitions mostly aligned with those found via our pol-

icy review (Section 3.2). Because we asked participants about platforms with divergent policies (i.e., Facebook and Twitter), we discuss each separately. Our final codebook is available in Appendix F. Figure 4 summarizes responses by platform. Because participants could describe multiple dimensions, these counts do not sum to the total number of participants. These numbers only represent front-of-mind definitions; not mentioning a dimension does not necessarily mean the participant does not believe that dimension applies to the policy.

Participants were more likely to believe Facebook confirms user identity. 54.2% of Facebook-assigned participants stated Facebook confirms the user’s identity matches their online persona. As one participant said, “[users] need to submit identification, and Facebook manually reviews it.” Only 25.1% of Twitter-assigned participants said the same. This difference was statistically significant ($\chi^2 = 140.58, p < 0.001$). While the share of Twitter-assigned participants who believe Twitter verifies identity is concerning, the majority of participants’ perceptions align with each platform’s actual policies. While not directly comparable, we note the percentage of participants stating Twitter verifies user identity in our survey is much lower than in Xiao et al.’s [78], potentially indicating user understanding of Twitter’s policy has improved.

Many participants focused on measures to ensure accounts were made by real humans, not bots. Instead of ensuring VAs’ true identity matched their persona, many perceived verification as simply requiring the user verify personal information (e.g., mailing address, email, phone number), limiting verification of bots (18.4% Twitter; 18.3% Facebook). We did not observe a statistically significant difference ($\chi^2 < 0.001, p = 1$). Both platforms require these checks, though they are Twitter’s primary verification mechanism.

Payment is mostly associated with Twitter. More than half of Twitter-assigned participants mentioned payment (56.0%). One participant explained, “You pay \$8 and elon gives you the blue checkmark.” Conversely, few (10.4%) Facebook-assigned participants believed payment was required. This difference was statistically significant ($\chi^2 = 370.6, p < .001$). Xiao et al. found similar results (i.e., Twitter is paid and Facebook is free) [78], but this would have been correct at the time, as their survey was conducted before Facebook switched to a paid model. We show this perception of Facebook as free has persisted and created a misconception among users, indicating they are unaware of Facebook’s policy change.

Facebook-assigned participants were more likely to believe verification was for “notable” accounts. 19.3% of Facebook-assigned participants said only notable accounts could be verified. As one participant said, “they have to be notable enough to where other people want to make fake accounts of them.” This misconception was not common, but was more common ($\chi^2 = 21.881, p < .001$) among Facebook-assigned than Twitter-assigned participants (10.8%). Con-

versely, Twitter-assigned participants (8.9%) were more likely ($\chi^2 = 44.80, p < 0.001$) than Facebook-assigned participants (1.4%) to say anyone could be verified.

Facebook-assigned participants were more likely to be unaware of the platform’s policy. Many Facebook participants reported not knowing Facebook’s policy (17.2%). One participant said, “I actually don’t know what the qualifications are to maintain a checkmark. I kind of blindly trust it has been adequately verified.” Some were even unaware Facebook had VAs (2.1%). One participant stated, “Facebook uses blue checkmarks? I thought you were talking about Twitter.” Many fewer Twitter-assigned participants (7.6%) reported lacking knowledge ($\chi^2 = 32.988, p < .001$).

Some people still conflate verification with credibility. Though not many, some participants (3.7% of Facebook-assigned; 2.7% of Twitter-assigned) continue to believe verification indicates the account is a reliable source of information. As one participant explained, “I would think that Facebook’s fact checkers would verify the post was legit and gave good information.” This mirrors previous work showing a minority of users conflate authenticity with credibility [12, 21, 72, 78].

Participants criticized Twitter more. Some participants mistrusted the verification process. They described it as politically biased (e.g., “They must share the same ‘opinion’ as Facebook’s creator/staff”), doing too little to prevent inauthentic accounts (e.g., “there are so many loopholes now for bots to act like humans and falsify information”), or expressed nihilism (e.g., “Better to let the [expletive] thing die than waste time on this verification nonsense”). Criticism was more common ($\chi^2 = 23.914, p < .001$) among Twitter-assigned (9.6%) than Facebook-assigned participants (3.4%). These are small fractions, but we note we prompted participants to share their definition of the process, not their opinion of it.

5.4 Verification Policy Perceptions (RQ3)

We next discuss perceptions of VA posts’ credibility after defining a verification policy (Part F) and how acceptable participants consider the policy (Part E). We saw a significant increase in perceptions that the VA’s posted content was credible ($Z = 21.69, p < 0.001$ in Wilcoxon Signed Rank test). This was likely affected by our priming participants to focus on verification by asking for a definition (Part D) and giving a specific policy (Part E). Therefore, we do not compare initial and after-priming responses, but only provide between-participant comparisons on the after-priming question.

We focus first on the three varied policy dimensions (Eligibility, Verification Method, and Payment), then discuss other factors. Figure 5a summarizes participant correctness perceptions, divided by dimension, and Figure 5b shows how acceptable participants considered each policy.

Limiting verification to notable accounts and authenticat-

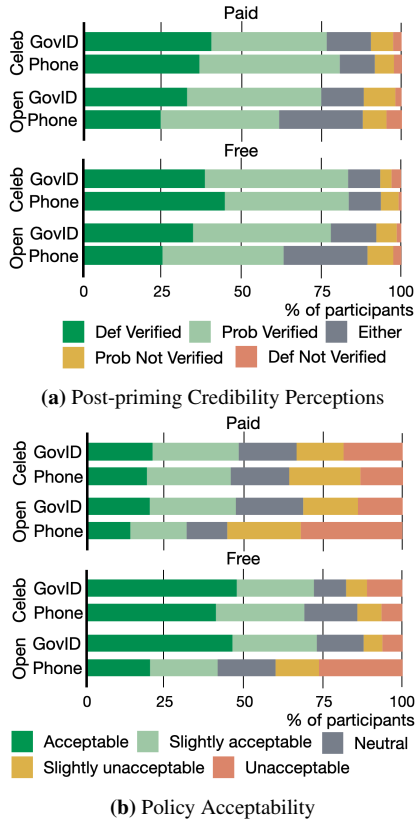


Figure 5: Likert-scale response indicating (a) posted content credibility perceptions and (b) policy acceptability after defining a policy. Both are organized by assigned policy dimensions.

ing using a government ID (govID) increases perceived posted content credibility and acceptability. The vast majority of govID-assigned participants (78.0%) believed the VA’s posted content was “Definitely” or “Probably” more credible. Perception of VA posted content credibility dropped to 72.0% when told accounts were verified via email or phone, with more participants indicating “Either” posted content could be credible (18.5%; 12.9% for govID). This difference was statistically significant, with govID-assigned participants 1.30× more likely to increase one point toward the VA ($p = 0.005$, Table 4b). GovID-assigned participants were also statistically significantly more likely to find the policy acceptable ($OR = 1.78, p < 0.001$, Table 5) with a majority finding it “Slightly acceptable” or “Acceptable” (60.0%), while this was a minority opinion for those shown the email or phone policy (46.9%). Requiring at least one form of govID was the most commonly desired policy change (N=306) with only 33 participants saying govID should not be required. One participant explained, “I would require a photo ID. I can say I’m John Travolta and I can give you my email address (which can be almost anything) to confirm me, but I’m not John Travolta.” This aligns with security best practices for verification [23], as it is much easier to create a new email or phone number than falsify a government document, and there have already

Variable	Value	Odds Ratio	CI	p-value
Eligibility	Anyone	–	–	–
	Notable	1.57	[1.32, 1.88]	<0.001*
Verification Method	Phone	–	–	–
	Gov ID	1.80	[1.51, 2.15]	<0.001*
Payment	Paid	–	–	–
	Free	2.53	[2.11, 3.03]	<0.001*
SA-6	1	–	–	–
	+1	1.27	[1.14, 1.41]	<0.001*

*Significant effect
– Base case (OR=1, by definition)

Table 5: Summary of regression over reported policy acceptability. Pseudo R^2 measures for the model were 0.04 (McFadden) and 0.11 (Nagelkerke).

been many cases of malicious accounts defeating phone verification [49, 63, 67].

There was a similar difference when comparing notable-only-assigned participants (80.8% “Definitely” or “Probably” more credible), as opposed to participants assigned an open policy (69.3% “Definitely” or “Probably” more credible). This difference was statistically significant with a slightly larger effect size ($OR = 1.80, p < 0.001$). Participants reported higher acceptability for the notable-only policy (58.5% “Slightly acceptable” or “Acceptable”), compared to an open policy (48.4% “Slightly acceptable” or “Acceptable”)—also statistically significant ($OR = 1.56, p < 0.001$). However, when asked for a desired policy change, a greater proportion of participants wanted the policy to be open, not notable. Of the 804 participants shown a notable-only policy, 18.9% wanted it to be open, while only 8.1% of open policy participants wanted verification for notable users only. This sentiment for open policies was driven by concerns of equality; as one participant stated, “I don’t believe one has to be well known or high-profile to be verified. That absolutely stinks of elitism.” This contradicts our regression results, suggesting participants are split on their preference for Eligibility.

Payment does not affect perceived correctness, but reduces approval. We did not observe a statistically significant impact on participants’ VA posted content credibility perceptions based on payment. When shown a free verification policy, 76.8% of participants indicated the VA’s post was “Definitely” or “Probably” more credible, compared to 73.2% of participants shown a paid policy. Free verification was the strongest factor increasing policy approval ($OR = 2.54, p < 0.001$). While 64.0% shown a free policy found it at least “Slightly acceptable”, only 43.1% said the same of paid policies. Like Xiao et al. [78], we found many participants focused on price when suggesting a policy change (N=342). One participant said, “Money shouldn’t be a barrier to doing public good.” This indicates payment might not impact users’ VA perceptions, but it displeases users, as observed with Twitter [29].

Social media use and security attitudes play a role. Par-

Participants who use Facebook were more likely to view the VA's posted content as credible (76.9% said "Definitely" or "Probably" more credible) compared to non-Facebook users (66.3% said "Definitely" or "Probably" more credible) ($OR = 1.54, p < 0.001$). Participants who reported taking more general security actions were more likely to view the VA's posted content as more credible ($OR = 1.21, p < 0.001$) and find the policy acceptable ($OR = 1.27, p < 0.001$). This may suggest the misconception that VAs are "secure," i.e., should be trusted over other accounts. However, prior work contradicts this [21, 72, 78], and few participants said verification indicates credibility (see Section 5.3). This may instead be an effect of the specific contrasting scenarios we chose, where the only major difference was the verified indicator and accounts were authoritative. Security-conscious participants may have been more likely to consider this difference.

6 Discussion

Our results reveal users' understanding of recent verification policy changes, along with their perceptions of the changes and other potential policies. We suggest social media platform verification policy improvements and discuss future work.

Many participants were aware of Twitter's transition to paid, open verification without a required identity check. While the results are not directly comparable, this seems to indicate improved user awareness relative to Xiao et al.'s earlier survey, which found many users believed Twitter performed rigorous identity checks [78]. Conversely, our participants were unaware of Facebook's policy changes, believing it remained free and restricted to notable accounts. This misunderstanding is not as consequential as incorrectly believing accounts undergo identity verification. However, our results suggest participants were more likely to perceive VA posts as credible when only notable accounts are verified, so this misunderstanding still introduces misplaced trust.

To improve trust in the verification process, platforms should employ rigorous ID checks. Participants were more likely to find the VA's posted content more credible when it was verified with a government ID, more likely to find government ID verification acceptable, and frequently suggested an ID check be added to improve verification. This shows users value identity verification over other requirements for bot prevention or account consistency. If Twitter transitions back to rigorous identity verification (as was rumored [51]), future work should consider whether perceptions of Twitter's policy improve, as we might expect from our hypothetical settings, or if these perceptions represent a one-way-ratchet and are already ingrained in the minds of users.

We also did not observe any statistically significant difference in the verified indicator's effect between platforms before priming participants about verification. When primed, participants shown Facebook's policy were statistically significantly more likely to find the VA more credible than those

shown Twitter's policy. This suggests users do not internalize these differences without priming, and because Facebook's policy is less well known, may default to their understanding from Twitter. As social media platforms change verification policies, they must educate users to avoid misunderstandings. This is especially important when changing government ID and notability requirements, as these significantly impacted perceived credibility, though future work must determine the best way to educate users.

Restrictions on account eligibility produced mixed results. Under a notable-only policy, participants were more likely to perceive the VA's posted content as more credible and find the policy acceptable. However, when asked to suggest changes to the platform, participants contradicted this sentiment by saying verification should be open to all users. One remedy suggested by a few participants ($N=19$) is a tiered approach to verification. As one participant suggested, "I think for public service accounts such as the fire department, police department, federal government, etc. there should be a more rigorous verification process." Similarly, some participants wanted the platform to evaluate users' authoritative credentials ($N=62$). This could include verifying hospital credentials of medical professionals or press credentials for journalists. Twitter somewhat employs this approach with special indicators for government and business accounts (🇺🇸, 🏢). Although users may prefer this in theory, prior work found users misunderstood both badges [78]. Future research should consider the impact of these indicators, especially in emergency situations when an account's authority is important (similar to our bomb threat examples) and under various Verification Method regimes to determine the interaction between these variables.

Perhaps the most polarizing verification change is switching to a paid model. Participants found paid policies unacceptable and wanted to remove payment, matching prior work [78]. However, we did not observe an effect from payment on participants' posted content credibility perceptions. We might have expected participants to be less likely to trust paying accounts, since Twitter's verified indicator has been described as a "scarlet letter" [29] and impostor accounts have been created [49]. However, it seems users correctly associate these problems with the lack of identity verification, not payment. This suggests that while payment might annoy users, it does not negatively impact how they evaluate VA posts.

Finally, participants were statistically significantly more likely to find the VA credible after priming about verification. This could be the result of asking participants to consider a hypothetical policy, but appears more likely due to priming effects. This could be problematic for platforms using policies that do not have rigorous identity verification. Malicious users may be able to fool others into believing their posts by drawing attention to their verified indicator. Future work should investigate situations where other information beyond the verified indicator varies between contradictory posts to measure the potential risk of social engineering attacks.

References

- [1] Census bureau data. <https://data.census.gov/>. (Accessed 08-30-2023).
- [2] Internet Archive. Internet archive: Wayback machine. <https://archive.org/web/>. (Accessed 09-10-2023).
- [3] Associated Press. Students criticize the University of North Carolina’s response to an active shooter emergency. <https://www.voanews.com/a/awash-in-social-media-how-are-us-police-learning-to-inform-the-public-better-after-shootings-7100938.html>, 2023. (Accessed 09-10-2023).
- [4] Brooke Auxier and Monica Anderson. Social media use in 2021. <https://www.pewresearch.org/internet/2021/04/07/social-media-use-in-2021/>, 2021. (Accessed 08-01-2019).
- [5] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300, 1995.
- [6] Cody Buntain, Jennifer Golbeck, Brooke Liu, and Gary LaFree. Evaluating public response to the boston marathon bombing and other acts of terrorism through twitter. *Proceedings of the International AAAI Conference on Web and Social Media*, 10(1):555–558, Aug. 2021.
- [7] Pew Research Center. Social media and news fact sheet. <https://www.pewresearch.org/journalism/fact-sheet/social-media-and-news-fact-sheet/>, 2022. (Accessed 08-18-2023).
- [8] Shelly Chaiken. The Heuristic Model of Persuasion. In *Social Influence: the Ontario Symposium*, volume 5, pages 3–39, 1987.
- [9] Brian X. Chen and Ryan Mac. Twitter’s blue check apocalypse is upon us. here’s what to know. *The New York Times*. <https://www.nytimes.com/2023/03/31/technology/personaltech/twitter-blue-check-musk.html>, 2023. (Accessed 09-10-2023).
- [10] Juliet Corbin, Anselm Strauss, and Anselm L Strauss. *Basics of qualitative research*. Sage, 2014.
- [11] H. Cramér. *Mathematical Methods of Statistics*. Princeton Landmarks in Mathematics and Physics. Princeton University Press, 1999.
- [12] Jazlyn Elizabeth Dumas and Rusty Allen Stough. When influencers are not very influential: The negative effects of social media verification. *Journal of Consumer Behaviour*, 21(3):614–624, 2022.
- [13] @elonmusk. Twitter’s current lords & peasants system for who has or doesn’t have a blue checkmark is bullshit. Power to the people! Blue for \$8/month. <https://twitter.com/elonmusk/status/1587498907336118274>, 2022. (Accessed 09-10-2023).
- [14] @elonmusk. Yes, this will destroy the bots. If a paid Blue account engages in spam/scam, that account will be suspended. <https://twitter.com/elonmusk/status/1587512669359292419>, 2022. (Accessed 09-10-2023).
- [15] @elonmusk. Given that modern AI can solve any “prove you’re not a robot” tests, it’s now trivial to spin up 100k human-like bots. <https://twitter.com/elonmusk/status/1640199090112806912?s=20>, 2023. (Accessed 09-10-2023).
- [16] Cori Faklaris, Laura A. Dabbish, and Jason I. Hong. A Self-Report measure of End-User security attitudes (SA-6). In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 61–77, Santa Clara, CA, August 2019. USENIX Association.
- [17] Andrew J Flanagan and Miriam J Metzger. The Role of Site Features, User Attributes, and Information Verification Behaviors on the Perceived Credibility of Web-Based Information. *New Media & Society*, 9(2):319–342, 2007.
- [18] B. J. Fogg, Cathy Soohoo, David R. Danielson, Leslie Marable, Julianne Stanford, and Ellen R. Tauber. How do users evaluate the credibility of web sites? a study with over 2,500 participants. In *Proceedings of the 2003 Conference on Designing for User Experiences, DUX ’03*, page 1–15, New York, NY, USA, 2003. Association for Computing Machinery.
- [19] Karl Pearson F.R.S. X. on the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *Philosophical Magazine*, 50(302):157–175, 1900.
- [20] Brian Fung. How Elon Musk transformed Twitter’s blue check from status symbol into a badge of shame. *CNN Business*. <https://www.cnn.com/2023/04/24/tech/musk-twitter-blue-check-mark/index.html>, 2023. (Accessed 09-10-2023).
- [21] Christine Geeng, Savanna Yee, and Franziska Roesner. Fake news on facebook and twitter: Investigating how people (don’t) investigate. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI ’20*, page 1–14, New York, NY, USA, 2020. Association for Computing Machinery.
- [22] Google. Verification Badges on Channels. <https://support.google.com/youtube/answer/3046484>. (Accessed 09-10-2023).
- [23] Paul A. Grassi, James L. Fenton, Naomi B. Lefkowitz, Jamie M. Danker, Yee-Yin Choong, Kristen K. Greene, and Mary F. Theofanos. NIST Special Publication 800-63A, Digital Identity Guidelines, Enrollment and Identity Proofing. Technical report, National Institute of Standards and Technology, 06 2017.
- [24] Amelia Hassoun, Ian Beacock, Sunny Consolvo, Beth Goldberg, Patrick Gage Kelley, and Daniel M. Russell. Practicing information sensibility: How gen z engages with online information. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI ’23*, New York, NY, USA, 2023. Association for Computing Machinery.
- [25] Andrew F Hayes and Klaus Krippendorff. Answering the call for a standard reliability measure for coding data. *Communication methods and measures*, 1(1):77–89, 2007.
- [26] Brian Hilligoss and Soo Young Rieh. Developing a Unifying Framework of Credibility Assessment: Construct, Heuristics, and Interaction in Context. *Information Processing & Management*, 44(4):1467–1484, 2008.
- [27] Kyle Hunt, Bairong Wang, and Jun Zhuang. Misinformation debunking and cross-platform information sharing through twitter during hurricanes harvey and irma: a case study on shelters and id checks. *Natural Hazards*, 103:861–883, 2020.
- [28] Christian Johnson and William Marcellino. Reining in COVID-19 Disinformation from China, Russia, and Elsewhere, November 2021. <https://www.rand.org/blog/2021/11/reining-in-covid-19-disinformation-from-china-russia.html>.
- [29] Alex Kirshner. How Elon Musk Turned the Blue Check Mark Into a Scarlet Letter. *Slate*. <https://slate.com/technology/2023/04/elon-musk-twitter-blue-check-marks-verification-lebron-james.html>, 2023. (Accessed 09-10-2023).
- [30] Ziva Kunda. The Case for Motivated Reasoning. *Psychological bulletin*, 108(3):480, 1990.
- [31] Ian Lamont. Plane Lands on the Hudson, and Twitter Documents It All. *Computerworld*. <https://www.computerworld.com/article/2530453/plane-lands-on-the-hudson--and-twitter-documents-it-all.html>, 2009. (Accessed 09-10-2023).
- [32] Annabelle Liang. Elon Musk: Twitter boss announces blue tick shake-up. *BBC News*. <https://www.bbc.com/news/business-65095684>, 2023. (Accessed 09-10-2023).
- [33] LinkedIn. Verified on your linkedin profile. <https://www.linkedin.com/help/linkedin/answer/a1359065>. (Accessed 09-10-2023).
- [34] Megan Loe. 5 VERIFIED Ways You Can Fact-check Online Claims. *Verify*. <https://www.verifythis.com/article/news/verify/fact-sheets-verify/5-tips-fact-check-online-claims-yourself-guide/536-64c58fc6-f17d-42dd-9970-b1b4814f9a87>, 2023. (Accessed 09-06-2023).

- [35] Gary Machado, Alexandre Alaphilippe, Roman Adamczyk, and Antoine Gregoire. Indian Chronicles: deep dive into a 15-year operation targeting the EU and UN to serve Indian interests. Technical report, EU Disinfo Lab.
- [36] Odanga Madung and Brian Obilo. Inside the Shadowy World of Disinformation-for-hire in Kenya. Technical report, Mozilla Foundation. Section: Fellowships & Awards.
- [37] Miriam Matthews, Katya Migacheva, and Ryan Andrew Brown. Superspreaders of Malign and Subversive Information on COVID-19: Russian and Chinese Efforts Targeting the United States. Technical report, RAND Corporation, April 2021.
- [38] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for csw and hci practice. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–23, 2019.
- [39] Adam W Meade and S Bartholomew Craig. Identifying careless responses in survey data. *Psychological methods*, 17(3):437, 2012.
- [40] Marcelo Mendoza, Barbara Poblete, and Carlos Castillo. Twitter under Crisis: Can We Trust What We RT? In *Proceedings of the First Workshop on Social Media Analytics*, SOMA '10, page 71–79, New York, NY, USA, 2010. Association for Computing Machinery.
- [41] Meta. Meta verified I get a verified blue check on instagram, facebook. <https://about.meta.com/technologies/meta-verified/>. (Accessed 09-10-2023).
- [42] Miriam J. Metzger, Andrew J. Flanagan, and Ryan B. Medders. Social and Heuristic Approaches to Credibility Evaluation Online. *Journal of Communication*, 60(3):413–439, 08 2010.
- [43] Meredith Ringel Morris, Scott Counts, Asta Roseway, Aaron Hoff, and Julia Schwarz. Tweeting is believing?: Understanding microblog credibility perceptions. In *Proceedings of the 15th ACM Conference on Computer Supported Cooperative Work*, CSCW '12, pages 441–450, New York, NY, USA, 2012. ACM.
- [44] Casey Newton and Zoe Schiffer. Elon Musk ignored Twitter's internal warnings about his paid verification scheme. *The Verge*. <https://www.theverge.com/2022/11/14/23459244/twitter-elon-musk-blue-verification-internal-warnings-ignored>, 2022. (Accessed 09-10-2023).
- [45] Matt Novak. Viral Video Alleging Canadian Wildfires Were 'Set Up' Is Very Misleading. *Forbes*. <https://www.forbes.com/sites/mattnovak/2023/06/09/viral-video-alleging-canadian-wildfires-were-set-up-is-very-misleading/?sh=67e194bb7350>, 2023. (Accessed 09-10-2023).
- [46] Matt O'Brien. Canada wildfire evacuees can't get news media on facebook and instagram. some find workarounds. *AP News*. <https://apnews.com/article/canada-wildfires-yellowknife-nwt-facebook-instagram-meta-723687efe632884e4eb1172528abb43f>, 2023. (Accessed 09-10-2023).
- [47] Matt O'Brien and Kathleen Foody. Confusion as Musk's Twitter yanks blue checks from agencies. *AP News*. <https://apnews.com/article/twitter-elon-musk-blue-checkmark-celebrities-544cfd66ed3a62f51a8a80c20e11ac5b>, 2023. (Accessed 09-10-2023).
- [48] Kari Paul. Russian disinformation surged on social media after invasion of Ukraine, Meta reports. *The Guardian*, April 2022. <https://www.theguardian.com/world/2022/apr/07/propaganda-social-media-surge-invasion-ukraine-meta-reports>.
- [49] Kari Paul. Fake accounts, chaos and few sign-ups: the first day of twitter blue was messy. *The Guardian*. <https://www.theguardian.com/technology/2023/apr/21/elon-musk-twitter-blue-rollout>, 2023. (Accessed 09-10-2023).
- [50] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. Beyond the turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70:153 – 163, 2017.
- [51] Sarah Perez. Twitter testing government ID-based verification, new screenshots show. *TechCrunch*. <https://techcrunch.com/2023/03/20/twitter-testing-government-id-based-verification-new-screenshots-show/>, 2023. (Accessed 08-18-2023).
- [52] Karena Phan. Social media videos push baseless conspiracy theory that blue items were spared from maui wildfires. *AP News*. <https://apnews.com/article/fact-check-conspiracy-blue-items-maui-wildfires-118319149774>, 2023. (Accessed 09-10-2023).
- [53] Prolific. Representative samples. <https://researcher-help.prolific.co/hc/en-gb/articles/360019236753-Representative-samples>, 2023. (Accessed 09-10-2023).
- [54] Adrian E Raftery. Bayesian model selection in social research. *Sociological methodology*, pages 111–163, 1995.
- [55] Harry T Reis and Charles M Judd. *Handbook of research methods in social and personality psychology*. Cambridge University Press, 2000.
- [56] David E. Sanger and Julian E. Barnes. U.S. Warns Russia, China and Iran Are Trying to Interfere in the Election. Democrats Say It's Far Worse. *The New York Times*, July 2020. <https://www.nytimes.com/2020/07/24/us/politics/election-interference-russia-china-iran.html>.
- [57] Craig Silverman. Verification and Fact Checking. *European Journalism Centre*. <https://datajournalism.com/read/handbook/verification-1/additional-materials/verification-and-fact-checking>. (Accessed 09-06-2023).
- [58] Craig Silverman and Rina Tsubaki. A guide to verifying digital content in emergencies. *Global Investigative Journalism Network*. <https://gijn.org/2014/03/18/a-guide-to-verifying-digital-content-for-emergency-coverage/>, 2014. (Accessed 09-06-2023).
- [59] Snapchat. How to verify your public profile. https://businesshelp.snapchat.com/s/article/public-profile-verify?language=en_US. (Accessed 09-10-2023).
- [60] Elliott Sober. Instrumentalism, parsimony, and the akaike framework. *Philosophy of Science*, 69(S3):S112–S123, 2002.
- [61] Truth Social. Red check verification. <https://help.truthsocial.com/moderation/how-to-get-verified>. (Accessed 09-10-2023).
- [62] Todd Spangler. Elon Musk Says Twitter 'Final Date' for Removing Legacy Blue Check-Marks Is 4/20. *Variety*. <https://variety.com/2023/digital/news/twitter-musk-date-removal-blue-checkmarks-legacy-1235570782/>, 2023. (Accessed 09-10-2023).
- [63] Mariana Spring and Laura Gozzi. Twitter blue tick: Multiple hillarys and new yorks as verifications disappear. *BBC News*. <https://www.bbc.com/news/technology-65346263>, 2023. (Accessed 09-10-2023).
- [64] Biz Stone. Not Playing Ball. *Twitter*. https://blog.twitter.com/official/en_us/a/2009/not-playing-ball.html, June 2009. (Accessed 07-18-2023).
- [65] Anselm Strauss and Juliet Corbin. *Basics of qualitative research*, volume 15. Newbury Park, CA: Sage, 1990.
- [66] S Shyam Sundar. The MAIN Model: A Heuristic Approach to Understanding Technology Effects on Credibility. *Digital media, youth, and credibility*, 73100, 2008.
- [67] Pete Syme. Elon musk's war against twitter bots isn't going very well. next, you'll have to pay to dm those who don't follow you. *Business Insider*. <https://www.businessinsider.com/elon-musk-war-on-twitter-bots-isnt-working-limits-dms-2023-6>, 2023. (Accessed 09-10-2023).
- [68] Jenny Tang, Eleanor Birrell, and Ada Lerner. Replication: How well do my results generalize now? the external validity of online privacy and security surveys. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 367–385, Boston, MA, August 2022. USENIX Association.
- [69] TikTok. Verified accounts on tiktok. <https://support.tiktok.com/en/using-tiktok/growing-your-audience/how-to-tell-if-an-account-is-verified-on-tiktok>. (Accessed 09-10-2023).

[70] Twitter. How To Get the Blue Checkmark on X. Twitter. <https://help.twitter.com/en/managing-your-account/about-twitter-verified-accounts>. (Accessed 09-10-2023).

[71] Twitter. Legacy verification policy. <https://help.twitter.com/en/managing-your-account/legacy-verification-policy>. (Accessed 09-10-2023).

[72] Tavish Vaidya, Daniel Votipka, Michelle L. Mazurek, and Micah Sherr. Does being verified make you more credible? account verification’s effect on tweet credibility. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI ’19, New York, NY, USA, 2019. Association for Computing Machinery.

[73] Sarah Vieweg. Microblogged contributions to the emergency arena: Discovery, interpretation and implications. *Computer Supported Collaborative Work*, pages 515–516, 2010.

[74] James Vincent. Twitter says paying blue subscribers now get ‘prioritized rankings in conversations’. <https://www.theverge.com/2022/12/23/23523845/twitter-blue-paying-priority-replies-conversations>, 2022. (Accessed 09-10-2023).

[75] Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics bulletin*, 1(6):80–83, 1945.

[76] Sam Wineburg, Sarah McGrew, Joel Breakstone, and Teresa Ortega. Evaluating Information: The Cornerstone of Civic Online Reasoning. *Stanford Digital Repository*, 2016.

[77] Stefan Wojcik and Adam Hughes. Sizing Up Twitter Users. Pew Research Center. <https://www.pewresearch.org/internet/2019/04/24/sizing-up-twitter-users/>, 2019. (Accessed 08-01-2019).

[78] Madelyne Xiao, Mona Wang, Anunay Kulshrestha, and Jonathan Mayer. Account verification on social media: User perceptions and paid enrollment. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 3099–3116, Anaheim, CA, August 2023. USENIX Association.

A Overview

In our appendices, we describe our web scraping process for policy collection (Section B), provide our survey text (Section C), and additional tables and figures not included in the main paper for brevity (Section D. The full set of mock posts shown to users in our survey, the full codebook of free-response questions, demographic questions, debrief text, and the timeline of policy changes we observed can be found at https://osf.io/a9y3j/?view_only=d2608dffe87f40c09885c4e55637ddeb.

B Policy Review Web Scraping Process

To capture each platform’s verification policies, we created a simple web scraper in Python using the BeautifulSoup4 and Selenium libraries. This script was run daily to pull each policy, compare it to the prior version, and record changes. Because we began our collection in February 2023, we used the Internet Archive’s Wayback Machine [2] to collect older changes to the platforms’ policies. Therefore, our review could be an under-approximation of changes in the period prior to our direct collection. However, we note that we were able to capture all major changes to Twitter reported in the news, and no other platform had major changes during this period. This process generated a dataset of timestamped verification policy changes for each platform.

C Survey Questionnaire

In this appendix, we provide the full text of our survey for one particular condition (Twitter post with police content with the verified indicator assigned to the declarative statement). Throughout, we provide heading indicating the section of the survey as shown in Figure 1. These headings were not included in the survey shown to participants and are only included here for readability.

----- Survey begins -----

(Consent, Part A)

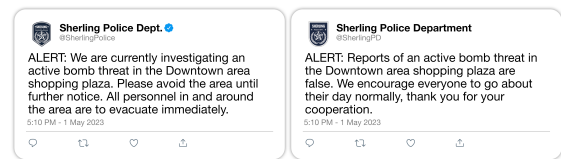
[Survey Consent presented here]

----- page break -----

In this study, we will display a pair of social media posts and ask you questions about the content shared in the posts.

----- page break -----

(Implicit Measure of Credibility, Part B)



Post A

Post B

Please answer the following questions considering the two Twitter posts above.

1. Post A and Post B contain conflicting information. Which of the posts do you believe is correct?
 - (a) Definitely A
 - (b) Probably A
 - (c) Equally likely to be A or B
 - (d) Probably B
 - (e) Definitely B

----- page break -----

(Explicit Measure of Credibility, Part C)

In this section, we will ask you some questions about how you determined which Twitter post was more correct in the previous section. Specifically, we will highlight different elements of the post and ask you how much each element influenced your decision. To help you know which visual element we’re asking about, we show a different Twitter post, distinct from the posts you saw before, and highlight the element in question.

A VA is denoted by a blue checkmark shown next to the display name, as illustrated within the red box below:



1. On the last page, we asked you which of two contradictory posts was more likely to be correct. When making that choice, how much did the presence of this verified account indicator (✓) affect your decision?

- (a) No Effect
- (b) Minor Effect
- (c) Moderate Effect
- (d) Major Effect

Every post on Twitter includes the display of the user’s profile picture next to their handle or username, as exemplified by the red box in the example post below:



1. On the last page, we asked you which of two contradictory posts was more likely to be correct. When making that choice, how much did the account’s *profile picture* affect your decision?

- (a) No Effect
- (b) Minor Effect
- (c) Moderate Effect
- (d) Major Effect

A display name is used to identify the account and can differ from the username. On Twitter, it appears next to the account’s profile picture as shown by the red box in the example post below:



1. On the last page, we asked you which of two contradictory posts was more likely to be correct. When making that choice, how much did the account’s *display name* affect your decision?

- (a) No Effect
- (b) Minor Effect
- (c) Moderate Effect
- (d) Major Effect

On Twitter, a user’s handle (also known as their username) is presented next to their profile picture on every tweet they post, and it is marked by the "@" symbol. An example of a user’s handle is provided in the red box below:

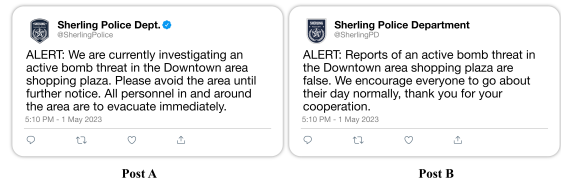


1. On the last page, we asked you which of two contradictory posts was more likely to be correct. When making that choice, how much did the account’s handle affect your decision?

- (a) No Effect
- (b) Minor Effect
- (c) Moderate Effect
- (d) Major Effect

----- page break -----

(Participants’ Definition of Verification, Part D)



One of the tweets you were previously shown was by an account with a verification checkmark (✓) indicating that the account has been verified.

1. Based on your understanding of Twitter’s account verification, what requirements must an account satisfy to become verified and obtain a verified checkmark?

----- page break -----

(Show Assigned Policy, Policy Acceptability, Suggested Changes, Part E)

Suppose Twitter adopted a verification policy in which the account had to meet all of the following criteria:

- **Any user** on the platform is allowed to apply for verification Accounts must submit government-issued identification that matches the name of the account being verified
- Any user on the platform is allowed to apply for verification Accounts must **submit government-issued identification** that matches the name of the account being verified
- Accounts **pay a monthly subscription fee** to maintain their verification checkmark

1. To what level do you believe these verification requirements are acceptable for verifying account owner identity?
 - (a) Unacceptable
 - (b) Slightly Unacceptable
 - (c) Neutral
 - (d) Slightly Acceptable
 - (e) Acceptable

2. If you could suggest one thing to add, remove, or change in this policy to improve its ability in verifying the account owner is who they say they are, what would it be? Please explain why.

----- page break -----

(Credibility Measure After Policy Priming, Part F)

We will now ask you to revisit the Twitter posts you were shown previously, and answer the following questions assuming this new policy was used for verification.

We display the Twitter posts and the new verification policy below for you to reference while you answer the questions.



- **Any user** on the platform is allowed to apply for verification Accounts must submit government-issued identification that matches the name of the account being verified
- Any user on the platform is allowed to apply for verification Accounts must **submit government-issued identification** that matches the name of the account being verified
- Accounts **pay a monthly subscription fee** to maintain their verification checkmark

1. Which of the following most closely resembles the subject matter of the two posts?
 - (a) Police investigating a bomb threat
 - (b) Effects of coffee on health
 - (c) Food recall due to E. coli outbreak

2. After reviewing the criteria required for an account to receive a verification checkmark, which of the posts do you believe is correct?
 - (a) Definitely A
 - (b) Probably A
 - (c) Equally likely to be A or B
 - (d) Probably B
 - (e) Definitely B
3. If a friend of yours was unsure about which post to trust, what would you say to this friend to help them decide?

----- page break -----

(Social Media Use, Part G)

Now we will end the survey with several short questions concerning your social media use and demographics.

1. Which of the following social media platforms do you currently have an account with? Select all that apply.
 - Twitter
 - Facebook
 - Instagram
 - LinkedIn
 - TikTok
 - YouTube
 - Other (please specify)
2. How often do you use Twitter in any given week?
 - (a) Daily
 - (b) Every other day
 - (c) Every two days
 - (d) Once a week
 - (e) I do not use Twitter
3. How often do you use Facebook in any given week?
 - (a) Daily
 - (b) Every other day
 - (c) Every two days
 - (d) Once a week
 - (e) I do not use Facebook
4. How much time do you spend on social media sites per day?
 - (a) Less than 30 minutes
 - (b) 30 minutes-1 hour
 - (c) 1-2 hours
 - (d) 2-4 hours
 - (e) 5-6 hours
 - (f) Greater than 6 hours

----- page break -----

(Security Attitudes, Part H)

Each statement below describes how a person might feel about the use of security measures. Examples of security measures are laptop or tablet passwords, spam email reporting tools, software updates, secure web browsers, fingerprint ID, and anti-virus software.

Please indicate the degree to which you agree or disagree with each statement. In each case, make your choice in terms of how you feel right now, not what you have felt in the past or would like to feel.

1. I seek out opportunities to learn about security measures that are relevant to me
 - (a) Strongly disagree
 - (b) Somewhat disagree
 - (c) Neither disagree nor agree
 - (d) Somewhat agree
 - (e) Strongly agree
2. I am extremely motivated to take all the steps needed to keep my online data and accounts safe.
 - (a) Strongly disagree
 - (b) Somewhat disagree
 - (c) Neither disagree nor agree
 - (d) Somewhat agree
 - (e) Strongly agree
3. Generally, I diligently follow a routine for security practices.
 - (a) Strongly disagree
 - (b) Somewhat disagree
 - (c) Neither disagree nor agree
 - (d) Somewhat agree
 - (e) Strongly agree
4. I often am interested in articles about security threats.
 - (a) Strongly disagree
 - (b) Somewhat disagree
 - (c) Neither disagree nor agree
 - (d) Somewhat agree
 - (e) Strongly agree
5. I always pay attention to experts' advice about the steps I need to take to keep my online data and accounts safe.
 - (a) Strongly disagree
 - (b) Somewhat disagree
 - (c) Neither disagree nor agree
 - (d) Somewhat agree
 - (e) Strongly agree
6. I am extremely knowledgeable about all the steps needed to keep my online data and accounts safe.
 - (a) Strongly disagree
 - (b) Somewhat disagree
 - (c) Neither disagree nor agree
 - (d) Somewhat agree
 - (e) Strongly agree

Factor	Description	Baseline
<i>Posted content Variables</i>		
Platform	The assigned visual design used to display posts	Twitter
Content type	The assigned content condition	Coffee
Position	The side of the argument the verified indicator was assigned to	Contradict.
<i>Policy Variables¹</i>		
Availability	Who can become verified?	Open
Verification Method	How are accounts verified?	Phone
Payment	Is payment required to become verified?	Paid
<i>Social Media Experience</i>		
Twitter experience	Does the participant report using Twitter (binary)	False
Facebook experience	Does the participant report using Facebook (binary)	False
Social Media Accts.	Number of social media platforms participants use	–
<i>Demographics</i>		
SA-6	Participant's score on Faklaris et al.'s SA-6 scale [16]	–
Age	Age of participant	–
Gender	Gender of participant	Male
Education	Does the participant hold a B.S. or higher degree (binary)	False

¹ Policy variables were only included when considering participants' policy acceptability rating (Part E) and their credibility perceptions after providing them with a mock policy (Part F).

Table 6: Factors used in regression models. Categorical variables are compared individually to the given baseline.

D Additional Tables and Figures

Finally, we provide tables and figures excluded from the main text for brevity. This includes a summary of the variables in the initial model for each regression (Table 6), additional participant demographics information (Table 7), and a summary of participants' responses regarding perceive impact of each account feature (Figure 6).

Metric	%	Metric	%
Gender		Income	
Woman	49.9%	<\$10k	10.6%
Man	48.4%	\$10k-\$25k	14.8%
Non-binary	1.2%	\$25k-\$50k	25.1%
Transgender/ Agender	0.3%	\$50k-\$75k	19.1%
Other	0.2%	\$75k-\$100k	11.9%
Race/Ethnicity		\$100k-\$150k	10.4%
White	73.9%	\$150k+	5.1%
Black	11.6%	Prefer not to respond	3.1%
Asian	6.0%		
Hispanic or Latino/a	4.9%		
Indigenous	0.7%		
Two or More Races	2.0%		
Other	0.2%		
Prefer not to respond	0.6%		

Table 7: Additional participant demographics.

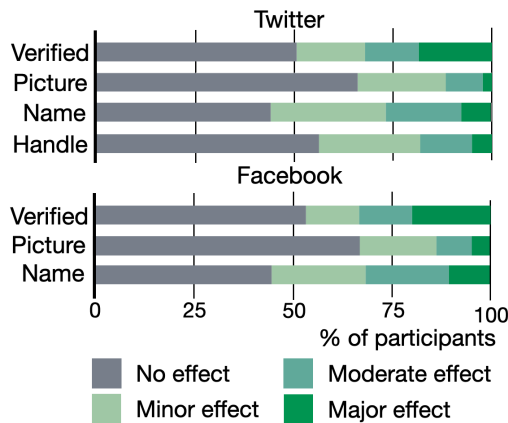


Figure 6: Likert-scale response indicating how much participants perceived each account feature impacted their credibility decision, organized by assigned social media platform.

E Demographics Questions & Debrief

(Demographics, Part I)

1. What is your age?
2. How do you describe your gender identity?
 - (a) Female
 - (b) Male
 - (c) Agender
 - (d) Non-binary
 - (e) Gender-queer
 - (f) Not sure
 - (g) Not listed above [with text entry]
 - (h) Prefer not to respond
3. Do you identify as Hispanic and/or Latino?
 - (a) Yes

- (b) No
- (c) Prefer not to respond

4. What level of education have you attained?
 - (a) Less than high school
 - (b) High School graduate (high school diploma or equivalent such as GED)
 - (c) Some college, but no degree
 - (d) Associate Degree
 - (e) Bachelor's Degree
 - (f) Master's Degree
 - (g) Professional Master's Degree (JD, MD)
 - (h) Doctorate Degree
 - (i) Prefer not to respond
5. What was your 2020 taxed income?
 - (a) Less than \$10,000
 - (b) \$10,000-\$24,999
 - (c) \$25,000-\$49,000
 - (d) \$50,000-\$74,999
 - (e) \$75,000-\$99,999
 - (f) \$100,000-\$149,000
 - (g) \$150,000 and greater
 - (h) Prefer not to respond
6. Do you get the majority of your earnings from Prolific or similar platforms?
 - (a) Yes
 - (b) No
 - (c) Prefer not to respond

----- page break -----

(Debrief, Part J)

Throughout this study you were shown social media posts showing conflicting reports about a particular event or research findings. These events are completely fictional and not based on any true events or findings. For the purpose of this study, these were made up to avoid bias in participant responses.

You were also given a set of criteria used for social media verification. Although the verification criteria we used for this study was based on the verification criteria Twitter and Facebook use to verify accounts on their platforms, the criteria you saw does not reflect the true criteria Twitter and Facebook use for their verification policies.

The verification process Twitter uses can be viewed in full by following this link. In this policy, verification is open to anyone but requires the owner of the account to pay a monthly fee to maintain the verification checkmark. The account must have a display name and profile photo. This display name and profile photo cannot be modified once the account has been verified. The account owner also must confirm a phone number with Twitter. Additionally, the account must show no signs of engaging in platform manipulation or spam, and show no signs of being misleading or deceptive.

The verification process Facebook uses can be viewed in full by following this link. This process is used for verifying accounts owned by public figures, celebrities, or notable brands. Notable brands are those that represent well-known, often searched for brands that are unique (i.e. be the only presence of this business), authentic (i.e. registered business), and have a complete Facebook Page or Facebook Profile (i.e. the account has a completed "About" section, has shared at least one post, and show recent activity).

Facebook also offers account profile verification for all accounts via Meta Verification. To be eligible for Meta Verification the account owner must be at least 18 years of age, have a public or private Facebook profile with the account owner's full name and a profile picture that matches a government issued ID. Additionally, the account must have a prior posting history, have two-factor authentication enabled. You can learn more about Meta Verification and its process here.

It can be difficult to determine whether information garnered online is true or false. However, there are steps you can take to help confirm if the information you read online is true or meant to mislead you. We provide links to several guides below for verifying digital content and fact checking information online below:

- [5 Ways You Can Fact-Check Online Claims](#)
- [A Guide to Verifying Digital Content in Emergencies](#)
- [Verification and Fact Checking - A General Guide](#)

F Codebook

In this appendix, we provide the codebooks used to categorize the valid free-text responses to the three open-ended questions in our survey (see Section C).

Table 8: Verification Definition

Code	Definition
<i>Identity Confirmation</i>	
Credentialed	The participant stated users must provide documentation of their credentials (e.g., medical license) to be verified.
Authentic	The participant stated that the verified indicator shows that the account actually belongs to the person indicated by the display name.
Government document verification	The participant stated a user must provide a copy of a government-issued document that verifies their identity.
Photo verification	The participant stated a user must submit a photo of themselves to verify their identity.
Banking verification	The participant indicated a person must provide their credit card number or banking information to verify their identity. This information is <i>not</i> used for payment.
Third-party verification	The participant stated a third-party service (such as id.me) verifies the user's identity.
<i>Consistent Persona</i>	
No recent changes	The participant stated that users cannot have made recent changes to their profile before receiving verification status.
Check other social media	The participant stated that Facebook/Twitter will check other social media accounts owned by the user in order to ensure the user is who they claim to be.
Knowledge verification	Users must answer questions that "only they would know" (i.e., security questions).
2FA enabled	An account must have two-factor authentication enabled in order to receive verification status.
IP address verification	Twitter/Facebook logs the user's IP address to ensure the verified account holder is the one logged in.
Posts match the account	The participant stated that Facebook/Twitter checks a verified account's post history to look for signs that the person posting has changed.
<i>Real Human</i>	
Email verification	The participant stated that Facebook/Twitter verified accounts by ensuring the account owner had access to their listed email address's inbox.
Phone verification	The participant stated a user must prove they own their provided phone number.
Location verification	The participant indicated a user must prove their address or physical location to Facebook/Twitter.
Only one account	The participant stated that the account owner cannot maintain other accounts on the same platform, unless the account is for another organization.
Long time user	The participant indicated that an account must be a specified number of days old in order to be verified.
Completed profile	The participant stated that the verified indicator meant that the Facebook/Twitter user had filled in all their profile information (bio, profile picture, etc.).
Active user	The account must be posting and/or interacting with other posts regularly for a minimum period of time.
Not a bot	The participant simply stated that Twitter guaranteed that the verified account was a real human, not a bot. However, they did not indicate that the human behind the account matched the displayed name.
<i>Eligibility</i>	
Notable	The participant stated that verified accounts were given to people of notoriety or public interest.
<i>Payment</i>	
Paid	The participant indicated monthly payment is required to be verified.
<i>Credibility</i>	

Reliable source of information	The participant stated that accounts with a verified indicator could be trusted to provide credible information.
Follows the rules	The participant indicated verification status is contingent on abiding by the terms and conditions of the platform, including no deceptive practices and no spamming.
<i>No Trust in Process</i>	
Political	The participant indicated that there was some political bias behind the accounts that received verification.
Verified accounts are dubious	The participant expressed distrust in verified accounts, or regarded them with suspicion.
Verification process is dubious	The participant believed the verification process is arbitrary, inconsistent, or insufficient to detect bad actors and impostors.
<i>Other</i>	
Visual indicator	The participant only described the visual appearance of the verified indicator, but not the process for achieving it.
Don't know	The participant did not provide a definition and simply stated that they did not know what was required to be verified.

Table 9: Policy Change

Code	Definition
<i>Identity Confirmation</i>	
Credentialed	The user should provide proof of any credentials (e.g., medical license) that they claim to have.
Government document verification	Users should be required to provide a copy of a government-issued document that verifies their identity.
No government document	Government-issued identification should <i>not</i> be required.
Photo verification	A user should be required to submit a “selfie” to verify their identity.
Banking verification	Users should be required to provide their credit card number or banking information to verify their identity. This information is <i>not</i> used for payment.
Employment verification	Facebook/Twitter should verify a user’s identity through the user’s employer.
Recorded video	Users should be required to upload recorded video footage of themselves to prove identity.
Interview	Users should be required to participate in an interview with the platform’s employees responsible for verification.
Criminal background check	All users must pass a criminal background check in order to receive verification status.
Multiple documents required	More than one government-issued identification document should be required for verification.
No anonymity	All verified users should be required to display their legal name with a clear photo of themselves.
Age verification	Facebook/Twitter should ensure each verified user is a legal adult.
Third-party verification	Facebook/Twitter should use a third-party service (such as id.me) to authenticate the user.
Biometrics	Users should be required to submit a fingerprint, facial scan, or retinal scan to verify the government-issued ID matches the account owner.
<i>Consistent Persona</i>	
No recent changes	Users should not have made recent changes to their profile before receiving verification status.
Check other social media	Facebook/Twitter should check other social media accounts owned by the user in order to ensure the user is who they claim to be.
Knowledge verification	Users must answer questions that “only they would know” (i.e., security questions).
2FA enabled	An account must have two-factor authentication enabled in order to receive verification status.

IP address verification	Twitter/Facebook should log the user's IP address to ensure the verified account holder is the one logged in.
Posts match the account	Facebook/Twitter should check a verified account's post history to ensure the same person has been posting.
Digital signature	Posts from verified accounts should be digitally signed.
Increase account security	The participant stated that verified accounts should have a higher level of security, though they did not specify what should be required to increase account security.
Recertification	Verified accounts should regularly undergo the verification process, or verification status will expire.
<hr/> <i>Real Human</i>	
Email verification	Facebook/Twitter should ensure the account owner has access to their listed email address's inbox.
No email verification	Email verification is an ineffective strategy.
Phone verification	Users should prove they own their provided phone number.
No phone verification	Phone verification is an ineffective strategy.
Location verification	Users should be required to prove their address or physical location to Facebook/Twitter.
Only one account	The account owner should not maintain other accounts on the same platform, unless the account is for another organization.
Long time user	An account should be open for a minimum period of time before it is eligible for verification.
Completed profile	The account should fill in all their profile information (bio, profile picture, etc.).
Active user	The account should be posting and/or interacting with other posts regularly for a minimum period of time prior to verification.
No active user requirement	The account should not need to be active to be verified.
Captchas	Accounts should be required to fill out a captcha before posting.
Not a bot	The participant simply stated that more must be done to remove bots from the platform.
<hr/> <i>Eligibility</i>	
Notable	Verification should be reserved for people of notoriety or public interest only. There should be separate tiers of verification, each with distinct verified indicators and different requirements for identity authentication. For example, government agencies would be required to submit more identifying documents than the average user, and the agency would have a different color verified indicator.
Tiered verification	
All accounts must be verified	Every account should be required to undergo the verification process.
Include anonymity	Users should be able to verify their identity to Facebook/Twitter without revealing their identity to other users.
<hr/> <i>Payment</i>	
Paid	Payment should be required to reduce impersonators and bad actors.
Reduced payment	The monthly payment should be reduced, or it should be a one-time fee.
No payment	Verification should be free to the user.
Free for public institutions	Government institutions should be exempt from verification fees.
<hr/> <i>Credibility</i>	
Follows the rules	Verification status should be contingent on abiding by the terms and conditions of the platform, including no deceptive practices and no spamming.
Label parody accounts	Parody accounts should contain some label to prevent users from believing they are conveying true information.
More fact checking	Facebook/Twitter must fact-check more posts.
Require explanation	Users should be required to explain why they deserve verification.
Transparent process	Facebook/Twitter should be more transparent about how they vet accounts and approve verification status for accounts.
Greater barrier to entry	The participant stated that it should be more difficult to achieve verification, though they did not provide an example of how to achieve this.
<hr/> <i>No Trust in Process</i>	

Skeptic	The participant was skeptical of verification in some way. These statements ranged from dubious (e.g., “The whole thing just needs to be revamped back to the way it was...[v]erification doesn’t mean anything anymore”) to abject fatalism (“There is nothing twitter can do at this point to return credibility to the checkmark system – it has been destroyed”).
<i>Other</i>	
No change	The participant did not wish to alter the presented policy.
Visual change	The participant only described a desired visual change the verified indicator.

Table 10: Friend Advice

Code	Definition
<i>Trust/Distrust Verification</i>	
Verified Account	The fact that the author was verified positively impacted the participant’s decision.
Unverified account	The participant distrusted the verified account because it was verified.
<i>Judge the content</i>	
Content	The participant made their decision based on the text of the post.
Check time posted	The participant recommended checking the time of the posts to discern whether one is more up-to-date. For example, if the post stating there is no bomb threat was more recent than the post alleging the bomb threat, perhaps the threat had been investigated and disproved.
Trust who cites sources	The participant recommended the friend trust whichever account cited a source, but did not indicate that the citation should be investigated.
<i>Judge the account</i>	
Vet the account	The participant recommended checking other elements of the account (e.g., name, bio, profile picture, and previous post history) to evaluate the account’s credibility.
<i>Trust the crowd</i>	
Check comments	The participant recommended checking the comments to see other users’ opinions of the posts.
<i>Trust other media</i>	
Other communication channel	The participant recommended contacting the same source via a different communication channel (e.g., call the police department directly).
Other sources	The participant recommended consulting other sources for information.
Trust neither	The participant cautioned against believing anything posted on social media.
<i>Trust yourself</i>	
Intuition	The participant recommended the friend trust their gut instinct or rely on their prior knowledge of the topic.
<i>Err towards caution</i>	
Better safe than sorry	The participant recommended choosing the safer option.