

Do Explanations Increase the Effectiveness of AI-Crowd Generated Fake News Warnings?

Ziv Epstein^{*1}, Nicolo Foppiani^{*2}, Sophie Hilgard^{*3}, Sanjana Sharma^{*4},
Elena Glassman³, David Rand⁵

¹ MIT Media Lab

² Physics Department, Harvard University

³ Computer Science Department, Harvard University

⁴ Graduate School of Design, Harvard University

⁵ MIT Sloan School

zive@mit.edu, ash798@g.harvard.edu, nicolofoppiani@g.harvard.edu,
ssharma@gsd.harvard.edu, eglassman@g.harvard.edu, drand@mit.edu

Abstract

Social media platforms are increasingly deploying complex interventions to help users detect false news. Labeling false news using techniques that combine crowd-sourcing with artificial intelligence (AI) offers a promising way to inform users about potentially low-quality information without censoring content, but also can be hard for users to understand. In this study, we examine how users respond in their sharing intentions to information they are provided about a hypothetical human-AI hybrid system. We ask i) if these warnings increase discernment in social media sharing intentions and ii) if explaining how the labeling system works can boost the effectiveness of the warnings. To do so, we conduct a study ($N = 1473$ Americans) in which participants indicated their likelihood of sharing content. Participants were randomly assigned to a control, a treatment where false content was labeled, or a treatment where the warning labels came with an explanation of how they were generated. We find clear evidence that both treatments increase sharing discernment, and directional evidence that explanations increase the warnings' effectiveness. Interestingly, we do not find that the explanations increase self-reported trust in the warning labels, although we do find some evidence that participants found the warnings with the explanations to be more informative. Together, these results have important implications for designing and deploying transparent misinformation warning labels, and AI-mediated systems more broadly.

Introduction

In recent years, there has been increased societal attention and governmental scrutiny on the spread of misinformation on social media platforms (Vosoughi, Roy, and Aral 2018), including concerns around political manipulation (Roose 2018) and public health (Bagherpour 2020) (e.g., the COVID-19 pandemic). In an effort to curb the spread of misinformation while avoiding direct censorship or content moderation, platforms such as Facebook and Twitter

have adopted policies that attach warning labels to certain posts, indicating that they may contain false or disputed information (Roth and Pickles 2020; Constone 2016; Mosseri 2017). Previous work has shown that labeling can reduce the rate at which users believe (Pennycook, Cannon, and Rand 2018; Pennycook et al. 2020; Clayton et al. 2019) and share (Yaqub et al. 2020; Pennycook, Cannon, and Rand 2018; Pennycook et al. 2020) misinformation.

However, these studies also highlight the enormous design space of such warning labels, whereby effectiveness could depend on factors such as the content and source of the label (Vraga and Bode 2017), the political affiliation of the user (Pennycook and Rand 2019), the presence or absence of warning labels on other items (Pennycook et al. 2020), label humorousness (Garrett and Poulsen 2019), and the signal word (Sherman, Redmiles, and Stokes 2020).

Results from explainable machine learning (Lai and Tan 2019; Bansal et al. 2020) suggest that adding explanations to a computer recommendation may increase the rate at which people accept that recommendation. While explanations have been shown to increase the trust and efficacy of human-AI systems in other domains (Lai and Tan 2019; Ribeiro, Singh, and Guestrin 2016) (particularly when the recommendations arise from a complex, algorithmic process (Rader, Cotter, and Cho 2018)), it is unclear how they will work in the context of social media, where attention is scarce, users are distracted by design (Pennycook et al. 2021b), and the key outcome is the effect on actual sharing and engagement behavior. In this work, we therefore ask whether explaining a system that generates misinformation labels may similarly increase user sharing discernment (the difference in sharing rates of true and false news items). To decouple the perceptions of warning labels with the particular accuracy of any given labeling system, we focus on a hypothetical system that is 100% accurate, i.e., only and all known misinformation receives a label. In particular, we study the role of explanations on warning labels generated by a (simulated) hybrid human-AI misinformation labeling system.

^{*}These authors contributed equally.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

We focus on a hybrid AI-crowd labeling system because it is highly practical relative to several alternatives. Many currently deployed misinformation labeling solutions rely on matching claims in social media posts to those disputed by professional fact checkers (Lyons 2021). However, these expert-based solutions cannot keep up with the scale and speed of misinformation, since misinformation can be produced faster than it can be disproven. Indeed, even among posts regarding known debunked claims, labeling can be inconsistent (Zannettou 2021). Due to these inherent delays, much violating content might never be labeled with a warning, and even those that do get labeled will only be after the period of maximal exposure. This inability of expert-based solutions to keep up with the volume of misinformation may also inadvertently encourage belief in misinformation through the “implied truth effect,” whereby users may perceive the lack of a warning label as validation that a piece of content has been fact checked (Pennycook et al. 2020). In addition, methods employing professional fact-checkers suffer from a general lack of bi-partisan trust: one study from the Pew Research Center found that 70% of Republicans and 48% of Americans say fact-checking efforts tend to “favor one side” (Walker and Gottfried 2019). Another proposed alternative is fully algorithmic methods for detecting misinformation (Bozarth and Budak 2020; Abdali et al. 2021; Weinzierl, Hopfer, and Harabagiu 2021) (see Related Work for a review). However, these methods fundamentally struggle to keep up with the non-stationarity of misinformation, and require ground-truth labeling. Hybrid AI-crowd systems represent a promising alternative that avoids the scalability problems of professional fact-checking by using non-expert human raters to crowdsource the evaluation of information accuracy (Kim et al. 2018; Pennycook and Rand 2019; Allen et al. 2020). These crowd ratings can then be used as inputs for machine learning algorithms, thereby combining the accuracy of human labels with the scale of automated systems. Platforms have already started to adopt such techniques, e.g. Facebook’s Community Review (Silverman 2019) or Twitter’s Birdwatch (Coleman 2021; Allen, Martel, and Rand 2021).

While such a system holds much promise, its complexity may undermine its effectiveness. Yaqub et al. (2020) suggest that warnings attributed to machine learning-based methods may be *less effective* than warnings attributed to human authorities in leading users to increase the quality of the news they share. They find that the warning effectiveness correlates with users’ self-reported “importance” of the sources, a measure they associate with trust in a given source, among other factors. However, the increased efficacy of fact checker and news media labels may also be at least partially attributable to a greater public familiarity with—and therefore greater perceived accuracy of—these label sources. Notably, Seo, Xiong, and Lee (2019) fail to replicate effectiveness of a fact-checker based warning with less than perfect accuracy.

With this in mind, we use the context of a hybrid crowdsourcing-machine learning system as we seek to understand the effectiveness of explaining warning labels in reducing intentions to share misinformation. Our inclusion of a baseline condition in which warning labels appear without

explanation additionally serves to help us understand users’ acceptance of such a system more generally. As Seo, Xiong, and Lee (2019) find explanations of a label attributed to machine learning increase user discernment without increasing self-reported user trust, we additionally investigate whether our effect is correlated with self-reported measures of trust. Concretely, we seek to further understand:

1. What is the baseline effectiveness of warning labels attributed to a hybrid crowdsourcing-machine learning system in reducing user intentions to share false news? Based on prior work (Seo, Xiong, and Lee 2019; Yaqub et al. 2020), we hypothesize the hybrid warning label will also be effective.
2. Is it possible to increase the effectiveness of labels attributed to a hybrid crowdsourcing-machine learning system by providing a detailed explanation of the system? Based on prior work (Seo, Xiong, and Lee 2019; Horne et al. 2019), we hypothesize explanations may increase sharing discernment.
3. Is increasing trust in the labeling system a necessary intermediary to increasing effectiveness of warning labels with respect to intention to share? We hypothesize that trust is not necessary to improve discernment (Seo, Xiong, and Lee 2019).
4. Do other moderators that have been investigated in prior work on misinformation interventions due to their connection to the psychology of fake news (Epstein et al. 2021; Pennycook and Rand 2021b,a) impact the effectiveness of providing a detailed explanation of the warning system?

We find clear evidence that warnings attributed to human-AI labeling systems increase sharing discernment relative to a no-warning control group. We also find some evidence that explaining the process of how this system works increases discernment above and beyond the warning itself. Furthermore, we find no difference in the measures of trust between the warning labels with and without explanation, suggesting that enhanced user trust is not the mechanism by which explanations may increase discernment. However, we find that certain moderators, such as the score in a cognitive reflection task and the age of the subjects, increase the efficacy of the explanations.

Related Work

Explanations and Transparency

The rise of interpretable or explainable machine learning is motivated, among other factors, by the idea that transparency in machine learning systems can build trust in their outputs (Lipton 2018; Doshi-Velez and Kim 2017). Research has shown that even nonsensical explanations (Lai and Tan 2019) or explanations of incorrect predictions (Bansal et al. 2020) can increase human acceptance of machine learning algorithms. Work on human explanations has shown that the effectiveness of uninformative explanations may be due to inattentiveness, as raising the stakes of a decision eliminates the effectiveness of uninformative explana-

tions, while retaining an effect of informative explanations (Langer, Blank, and Chanowitz 1978).

Kizilcec (2016) studied the effect of varying levels of explanation detail on trust in an algorithmic grading system and found that a medium level of transparency—in which an explanation of the grading *procedure* was provided—tended to increase trust beyond the low transparency condition, while providing specific details of the calculation (high transparency) decreased trust. Kizilcec further finds that the decrease in trust is correlated with a violation of expectations in the output of the algorithm. Kulesza et al. (2013) used a music recommendation environment to study the effects of varying the soundness and completeness of explanations on users' trust in the system. They find that high soundness and completeness together result in high trust in the system, but that users prefer concrete feature-based explanations to explanations of the algorithmic procedure.

Rader, Cotter, and Cho (2018) explore the effects of different types of explanations (*What, How, Why* and *Objective*) on users' perceptions of the Facebook news feed algorithm. They find that while these explanations generally increase knowledge and awareness of the algorithmic system, they may decrease user assessments of the system's correctness, fairness, and responsiveness to user controls. Liao, Gruen, and Miller (2020) developed an Explainable AI question bank to explore the distinct affordances of different types of explanations. They found that how explanations were helpful when users were in a "quality-control role."

Explanation depth is also studied in the context of debunking low quality information. Chan et al. (2017) present a meta-analysis of debunking studies, in which detailed debunking messages are found to be more effective than less detailed messages. The authors hypothesize, in line with previous works, that a debunking message must be sufficiently detailed to provide users with an alternative account of the false phenomenon. However, Lewandowsky et al. (2012) recommend using the simplest messages that satisfy this property. In line with this, Martel, Mosleh, and Rand (2021) show that correction style (including simple versus deep explanations) has minimal impact on correction efficacy.

In the context of AI-generated flags for misinformation on social media, Horne et al. (2019) compared warning labels with explanations to those without in an experimental design similar to ours. They however focus on article reliability as a key outcome, while we focus on sharing. They find that warning labels generated with AI improve perceptions of reliability, and find heterogeneity among participants' baseline perceptions of reliability.

Approaches to Detecting Misinformation

In an attempt to detect misinformation at a scale necessary for social media, several approaches have been proposed. One class involves computational methods to detect misinformation (Bozarth and Budak 2020). These computer-based approaches have shown remarkable success, leveraging signals such as sharing patterns (Rosenfeld, Szanto, and Parkes 2020), text features (Granik and Mesyura 2017), account activity (Breuer, Eilat, and Weinsberg 2020), user stance (Weinzierl, Hopfer, and Harabagiu 2021) and vi-

sual features for website screenshots (Abdali et al. 2021). However, the nuanced nature of truth, the limited availability of labeled training data (Rubin, Chen, and Conroy 2015; Bozarth, Saraf, and Budak 2020), and the non-stationarity problem whereby the signatures of misinformation can change rapidly (e.g. with the rise of COVID-19 misinformation) place fundamental limits of the effectiveness of fully automated systems.

A promising alternative that avoids these concerns involves crowdsourcing the fact-checking process. Indeed, crowdsourced fact checking could demote platforms as centralized authorities of truth, and keep up with the fast pace of new forms of misinformation. Kim et al. (2018) develop an algorithm for deciding which crowd-flagged items to send to professional fact-checkers for verification, with the goal of minimizing overall exposure to false news. Using datasets from Weibo and Twitter, they demonstrate that this method can be effective in reducing the spread of misinformation.

Pennycook and Rand (2019) show that laypeople can successfully distinguish high quality media outlets from those which produce hyper-partisan and blatantly false news. They also show layperson trust judgments of websites are highly correlated with those of professional fact-checkers and propose a modification to the newsfeed algorithm, whereby content from websites which receive low crowdsourced trust judgments is downranked. Epstein, Pennycook, and Rand (2020) build on this work by showing that when informed that their judgments will be used to fight misinformation on social media, participants still correctly distinguish between high and low quality outlets. This suggests that users would not game such a crowdsourcing mechanism to advance their partisan agendas. Allen et al. (2020) extend this work to the evaluation of headlines, instead of websites. They show that the ratings of small politically-balanced crowds are equally correlated with fact-checker ratings as the fact-checkers' ratings are correlated with one another.

Labeling Misinformation and Credibility Indicators

The research area of misinformation labeling is rapidly evolving, as academic and industry researchers are eager to understand the full implications of design decisions being made on social media platforms (Morrow et al. 2020). Clayton et al. (2019) investigate the effects of the locality and severity of misinformation warnings. They evaluate conditions including a general misinformation warning and advice on how to spot misinformation, with or without specific misinformation labels appealing to fact checkers. Further, fact checker labels vary between describing headlines as "rated false" and "disputed." They find that specific labels are more effective than general warnings at reducing the perceived accuracy of false news, and that general warnings have a spillover effect, also reducing the perceived accuracy of true news. In the Appendix, the authors additionally state results for intention to share, finding that users in the labeled conditions are no less likely to share items labeled as "rated false" compared to unlabeled items.

Pennycook, Cannon, and Rand (2018) find that prior exposure to false news, even when labeled as "Disputed by 3rd Party Fact-Checkers," increases perceived accuracy of

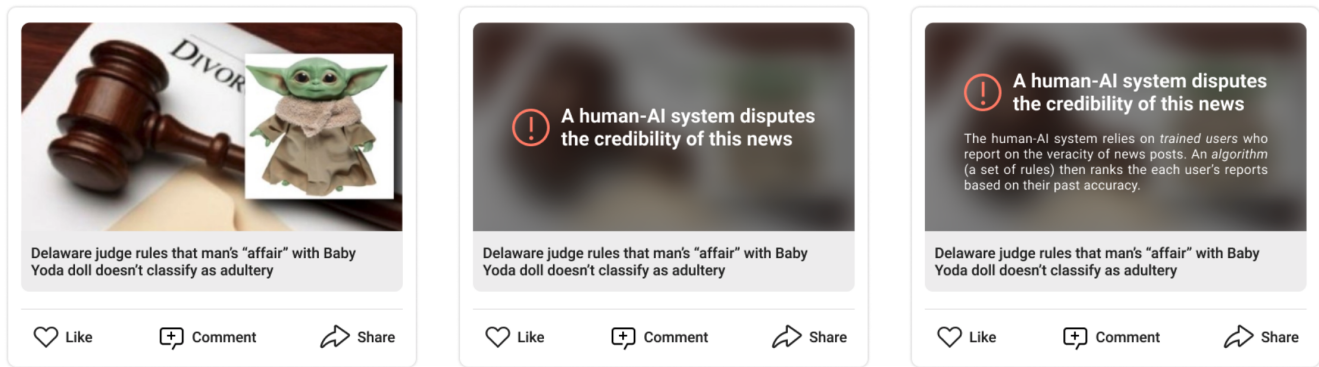


Figure 1: Example of a false post across the three study conditions: *the Control condition* (far left), *the Label condition* (center), and *the Explanation Label condition* (far right).

those news items when later viewed without labels. While the primary outcome of interest was perceived accuracy, the authors identify in a secondary analysis that disputed labels decrease users' intentions to share false news items. Pennycook et al. (2020) study both perceived accuracy and sharing intentions in the presence of labels and find that the presence of "disputed" or "false" labels decreases both accuracy and sharing for labeled false items, but reverses for unlabeled false items, via the "implied truth effect." The authors find that this effect can be mitigated by explicitly labeling some verified items as "true."

While early research in this field often focused on accuracy perceptions as outcome (Clayton et al. 2019; Pennycook, Cannon, and Rand 2018; Bode and Vraga 2015; Horne et al. 2019), recent work has shown that accuracy perceptions are not a reliable proxy for intention to share (Pennycook et al. 2021b), and thus addressing the goal of reducing the spread of misinformation requires measuring the latter.

Yaqub et al. (2020) explore the effects of appealing to different authorities in social media misinformation labels and find that users are most likely to change their sharing behavior in response to warning labels citing "fact checkers" or "news media" as the source of the dispute. A condition citing "AI" was among the least effective at reducing misinformation sharing in the study. However, this may be due to the fact that currently deployed misinformation labels generally appeal to fact checkers, and so users currently have a better understanding of the accuracy they can expect from fact checkers. It is difficult to know how users with no past experience would have assigned importance to equally accurate label sources. Seo, Xiong, and Lee (2019) similarly find some evidence that users are more likely to correctly detect false news when it is identified by "fact checkers" rather than "machine learning." However, this result did not replicate in a follow up study in which they removed news sources. This follow up study also tested the effect of adding more (fabricated) detail to the "machine learning" warning label and found that users were better able to detect false news in this condition but displayed less self-reported trust in the system, which may be due to the explanations being incongruous with the headlines.

Methods

We recruited $N = 2512$ Americans using Lucid, which uses quota matching to provide a sample that is nationally representative on age, gender, ethnicity and geographic region (Coppock and McClellan 2019). $N = 1855$ of these individuals passed two initial trivial quality checks at the start of the survey and were allowed to proceed. Additionally, we only considered subjects who self-reported using at least one social network among Facebook, Twitter, and Instagram. Of these participants, $N = 1473$ completed the discernment task, and $N = 1394$ completed the entire survey. Unless otherwise specified, we will consider the $N = 1473$ subjects who completed the discernment task as our final sample. This final sample had a mean age of 47.87, was 54.1% female, and was 87% White. Participants were shown 20 social media "cards" during the main discernment task, half with a false headline, half with a true headline. Each participant saw the 20 cards in a random order, and rated their sharing intentions ("How likely would you be to share this story on social media?") on a Likert scale from 1 (Extremely Unlikely) to 6 (Extremely Likely). These headlines were selected from a larger set of headlines (Pennycook et al. 2021a) to be relevant to news when the study ran, (e.g. "stale" headlines were pulled). The stimuli can be found online at https://osf.io/tgwna/?view_only=efa0ae9ebb30408ca54204e0dc3f53b6.

In a between subjects design, participants were assigned to one of three conditions: i) a control condition with no labeling (*the Control condition*); ii) a basic labeling condition in which users are shown a statement about how the labeling system works prior to any sharing decisions and then presented with the 20 headlines in which false headlines contain a simple label indicating that the claim has been labeled as false by the human-AI system (*the Label condition*); and iii) an explanation condition which contains all elements of the basic flagging conditions and augments each label with a general (static) explanation of the process by which that label was generated (*the Explanation Label condition*). See Figure 1 for examples of the false items across conditions.

After the discernment task, all subjects were asked two factual manipulation checks about what they had seen in

	Estimate	Standard Error	p (bayes)	Rhat
veracity	0.267	0.106	0.008**	1.002
Label	-0.155	0.094	0.052	1.007
Explanation Label	-0.161	0.048	0.031*	1.003
veracity × Label	0.135	0.064	0.019*	1.003
veracity × Explanation Label	0.217	0	< 0.001***	1.002
Intercept	3.201	0.099	0***	1.003
N=29,460				

Table 1: Bayesian mixed effects model predicting sharing intent. P (bayes) is the number of posterior samples greater than (less than) zero. Rhat measures the potential scale reduction factor (PSRF, (Brooks and Gelman 1998)), which captures the variation within a given chain relative to the variation across chains (at convergence, Rhat = 1).




Figure 2: Posterior distributions of the main effect of veracity (e.g. discernment in control, red), and the interactions between veracity and *the Label condition* and *the Explanation Label condition*, plotted using kernel density estimations. The blue and green distributions correspond to the increase in discernment above and beyond the control for the *the Label condition* and *the Explanation Label condition*, respectively. The coloured points show the average values of the distributions, while the error bars illustrate 95% credible intervals.

the images used during the main task (Kane and Barabas 2019). These questions ensure the treatment was deployed successfully. The subjects who received warning labels then answered a battery of questions related to their opinions of and preferences for the labeling process. Subjects answered a series of multiple choice questions related to how helpful, annoying, and informative these labels are, as well as how often the user would like to see them on social media. Finally, they answered a battery of six trust questions derived from Mayer, Davis, and Schoorman (1995)’s three factors of trustworthiness: ability, benevolence and integrity (ABI). For these six ABI questions, we find they are highly correlated ($\alpha = 0.821$), so following standard practice we use Principal Component analysis to construct a single measure of trust that explains 65.3% of the overall variance.

Users were then asked additional measures, such as two measures of attention (Berinsky et al. 2021), a 6-question cognitive reflection test (CRT) that measures the tendency to stop and think versus going with your gut (Frederick 2005; Thomson and Oppenheimer 2016), political affiliation, and basic demographics. These covariates were chosen because prior work has looked at how they moderate discernment (Epstein et al. 2021; Pennycook and Rand 2021a,b), so here we investigate if they impact the effectiveness of the explanation. Reaction times were not recorded.

For our main analysis, we fit a Bayesian mixed effects linear model predicting sharing rate as a function of con-

dition and headline veracity and their interaction. Our key tests are on the interactions between veracity and the treatment dummies (testing whether each treatment increases discernment – the difference in sharing of true relative to false news – relative to control). The data is analyzed at the response level, while mixed effects account for variability across subjects and headlines. We adopt a Bayesian modelling approach because the multilevel models with our mixed-effects structure did not converge. We perform full Bayesian inference on the model parameters using the R package BRMS (Bürkner 2017), which relies on Stan (Carpenter et al. 2017) to sample the posterior distribution using Markov Chain Monte Carlo. We assigned vague and weakly-informative priors, which allow the data to “speak for themselves” (the model specifications are outlined in the Section 3 of the Appendix). The significance of the coefficients extracted from the data is used as a metric to quantify the observation of an effect. For a given coefficient, we quantify significance with a “Bayesian p-value” (pbayes) by counting the percent of posterior samples that are below (or above) zero. The experimental design and analysis methods were preregistered before collecting the data, and are available at <https://aspredicted.org/m3dc7.pdf>.

Results

The results of our main analysis are shown in Table ???. We observe a significant interaction between *the Label condition*

and veracity, indicating that even the simple warning label of the human-AI labeling increases discernment (the difference in sharing of true and false news) relative to the control ($p_{\text{bayes}} = 0.019$, preregistered). We also observe a significant interaction between the *Explanation Label condition* and veracity, such that explaining the label also increases sharing discernment ($p_{\text{bayes}} < 0.001$, preregistered).

This main result is summarised in Figure 2, where the marginal posterior distributions for baseline discernment (i.e., in control) is plotted, together with the two marginal posterior distributions representing the increase in discernment for the warning label and the warning label with explanation conditions, respectively.

As can be seen from the two interaction coefficients in Table ??, adding the explanation roughly doubled the effect of the warnings on sharing discernment. To assess whether this increase from explanations above and beyond the minimal warning label was statistically significant, we compute the posterior distribution of the difference between the two interaction terms. The mean value is positive, reflecting the observed increase of discernment, but not at a significant level ($p_{\text{bayes}} = 0.110$, pre-registered). This result suggests that an explanation may increase the efficacy of the label above and beyond the warning label itself, but a conclusive answer would require an additional, larger study.

The Role of Trust

Next, we look at how perceived trust varies across conditions. To do so, we perform a linear regression predicting trust, defined as the first principal component of the ABI questions, as a function of the experimental condition (dummy) and controlling for attentiveness. We do not observe any significant difference in trust between the *Label condition* and the *Explanation Label condition* ($p=0.3293$). In fact, we directionally observe that more attentive subjects in the *Explanation Label condition* report lower trust ($p=0.09$) but the overall moderation of attentiveness is not significant either.

This suggests that, in so much as the explanation may increase the effectiveness of the warnings, trust in the warning label is not the mechanism by which this occurs. In addition, we find no significant differences between the two *Warning Label* conditions in participant ratings of how helpful the warning label was ($p=0.253$), how annoying the warning label was ($p=0.345$), how well the warning label helped the user understand the labeling process ($p=0.169$) or if the warning labels helped them understand what signals were leveraged ($p=0.56$). We do find that participants in the *Explanation Label condition* said the warning labels provided new information marginally more than in the *Label condition* ($p=0.067$). This may suggest that the explanation provides new information into the process that generated the the labels that the user can use or act on in their sharing decisions.

Factual Manipulation Check

After the main discernment task, we included manipulation checks to see what participants thought they saw during the discernment task. These results of these manipulation

checks is shown in Figure 3. Several subjects in the *Control condition* report to have seen a warning label, and some participants in both the *Label condition* and the *Explanation Label condition* stated they did not (see the left pane of Figure 3). Participants who said they saw a headline tagged with a warning message were then asked which warning labels they saw. While participants in the *Label condition* made up the plurality of respondents who saw “some headlines were labeled false,” they were also more likely to say they saw a warning label with an explanation than one without (see the right pane of Figure 3). This suggests that our warning label design worked up to a certain extent, but that even in a survey context — where people are likely paying more attention than when scrolling through social media feeds — users do not always realise they are presented with warning labels or explanations, or they might believe they have seen warning labels when they were not presented with them.

Additional Moderators

Additional analyses help us determine which factors moderate the treatment effect. We consider attention, propensity to go with one’s gut, as measured by the critical reflection test (CRT), partisanship, age, education and gender as possible moderators. Partisanship is coded as increasing means more Republican. Gender is coded as 1 as female and 0 otherwise.

For each of these factors, we modify the model for the main analysis by including this factor as a predictor, and re-run the Bayesian model. Each moderator is scaled by subtracting the average and dividing by the standard deviation. This ensures that the coefficients for the fixed effects are compatible with the main analysis within statistical uncertainties.

We observe that more attentive participants, higher CRT participants, more conservative participants, older participants, less educated participants, and female participants share less at baseline. More attentive participants, higher CRT participants and older participants are more discerning in their baseline sharing. However, we do not observe any significant interactions involving moderators, except for Education interacting with the *Label condition*, and attention interacting with the *Explanation Label condition* (We also do not account for multiple comparisons here, which would render both interactions non-significant). For statistical details, see Supplemental Table S1. This general lack of moderation suggests that warning labels work equally for a diverse set of users. However, this result fails to give us a sense of the mechanisms by which warning labels increase discernment, which should be explored in future work.

To better understand the mechanisms behind the increase in discernment due to the explanation itself, we also measure and report the difference in moderation of the treatment effect on discernment between the *Explanation Label condition* and the *Label condition* (see Figure 4) By calculating a credible interval for the difference between the two coefficients, we can assess if explanations are differentially effective for certain subgroups. We find a significant difference between the 3-way interactions for age, CRT, partisanship and education. This suggests that the explanations are differentially more effective at increasing discernment among




Figure 3: Factual manipulation check for participants in the two treatment conditions.

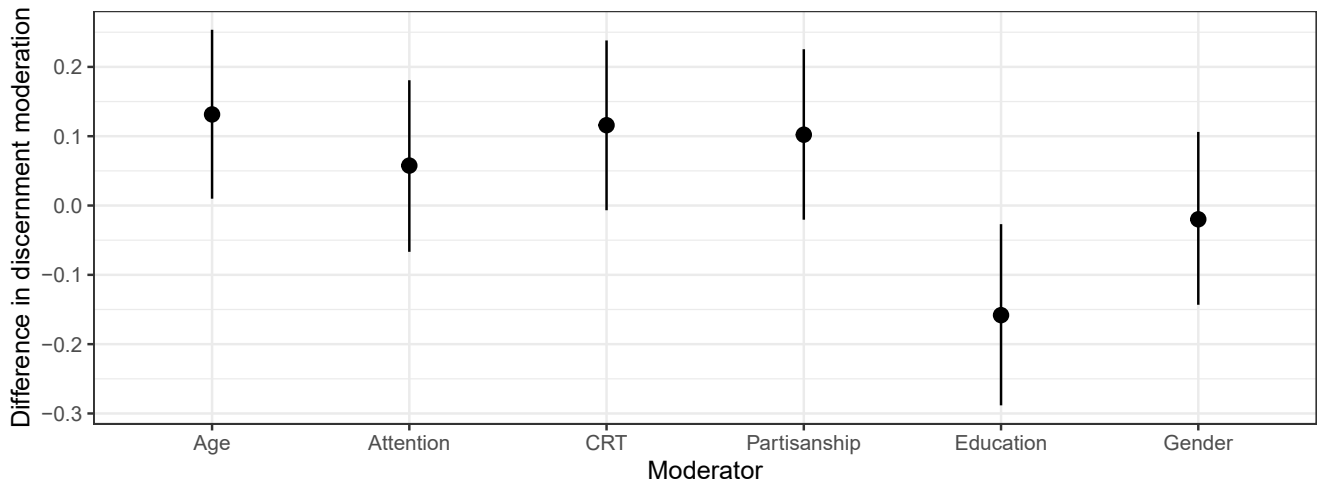


Figure 4: Posterior distributions for the difference in moderation of discernment between the two treatment conditions.

participants with lower levels of education or CRT, who are older, and are more conservative. We also observe a directional difference in 3-way interactions for attention, whereby explanations are more effective for attentive users.

Discussion

Our results suggest that providing users with warning labels created by a hybrid AI-crowd system can increase sharing discernment relative to a no-label control group. This suggests that labeling with human-AI warning labels has promise. Given the pragmatic merits of this approach for detecting misinformation at scale, our work suggests that people may indeed respond well to such a labeling scheme.

We also found circumstantial evidence that providing more transparency via explaining what the AI system is doing may boost the system’s efficacy above and beyond the basic warning. This adds to the conversation about correction explanations for debunking. Most work on corrective explanations focuses on explanation depth at the headline or

claim level (Lewandowsky et al. 2021), but has not looked at the corrective depth of the debunking *mechanism* itself. Ultimately, this approach is more scalable since all headlines share the same mechanism explanation, and generating headline-level explanations can be costly.

We also observed several significant moderators: explaining the warning labels was more effective for participants with lower levels of education or CRT, who are older, and are more conservative. These findings have important implications for how future explanations are designed, and at whom they are targeted. We also found directional moderation of attention, whereby explaining the warning label was more effective for more attentive users. Even in the survey context, where participants are presumably focused on the task at hand, users did not always realize they were presented with warning labels or explanations in the first place, or alternatively they might believe they saw labels when they were not present. While explanations may provide an important tool for boosting the efficacy of opaque human-AI systems, this effectiveness could be undermined by envi-

ronments where users are distracted and thus do not even consider them. This finding suggests important avenues for future work. First, it highlights the importance of designing warning labels in general, and explanations in particular, that grab more attention from users.

Further, our results suggest that (self-reported) user trust in the warning label is not the mechanism by which explanations contribute to the efficacy of machine recommendations. Despite a good deal of research that assumes a causal relationship between interpretability, trust, and acceptance (Ribeiro, Singh, and Guestrin 2016), previous works have been equivocal in proving these relationships. Cramer et al. (2008) find that while explanations increase acceptance of machine recommendations, they do not increase self-reported trust or willingness to delegate to the system on a future task. Papenmeier, Englebienne, and Seifert (2019) find that while explanations do not increase self-reported trust, they do increase *observed* trust, as measured by the probability that users switch their pre-reported answers to match those of the machine recommendation. In their experiment, this does not necessarily increase accuracy, as the algorithm does not achieve 100% accuracy. Indeed, we note that while achieving greater acceptance of machine recommendations is equivalent to increasing discernment in our experiment, in general the goal of explanations should be to increase *calibrated trust*, that is, trust that is appropriate for the performance of the system (Zhang, Liao, and Bellamy 2020). Papenmeier, Englebienne, and Seifert (2019) additionally find that the accuracy of the machine learning system has a greater effect on user trust, both self-reported and observed, than the presence and quality of explanations.

Our results are consistent with (Horne et al. 2019). While they consider accuracy instead of sharing as the key outcome, our work corroborates the idea that providing explanations to AI labels has potential to augment user discernment. Further, our results support their findings on the heterogeneity of the efficacy of such explanations. Our results concerning attentiveness and explanation effect are also consistent with Langer, Blank, and Chanowitz (1978). Studying human explanations (rather than AI), the authors found that relevant explanations are in low-stakes problems no more effective than a placebo explanation at increasing user acceptance of a request. However, when the stakes were slightly higher, the quality of the explanation provided did matter. The authors hypothesize that the higher stakes of the request induce attentiveness where respondents were previously inattentive, thereby increasing the importance of quality explanations. We believe this suggests interesting future directions for better understanding the moderators of warning label effectiveness, including evaluating users' trust with more measures than explicit self-reporting and understanding users' pre-existing experience with a given label source (e.g., fact checkers, AI, crowdsourcing) and perceptions concerning the accuracy of that label source.

Our work has several limitations that must be considered. First, our work only considers hypothetical sharing intentions, and therefore may suffer from demand effects. Since previous work has shown that self-reported sharing intentions correlates with actual sharing on Twitter (Mosleh, Pen-

nycook, and Rand 2020) and has replicated misinformation reducing interventions in Twitter field experiments (Pennycook et al. 2021b), we believe the use of online surveys to be justified by providing a sandbox to test otherwise challenging interventions. Second, in our study only false headlines were labeled by the system. Future work should explore how discernment changes when the system erroneously labels true items. Third, the nature of the human-AI explanation we used, with a vetting process for users and a particular crowd-sourcing methodology, is a specific representation from a large space of possible systems. We chose this particular process because of its similarity to real world systems, but future work should explore how these results generalize to across other possible systems, such as different kinds of explanations, differently designed labels, whether the AI labels or vets, and expert versus non-expert crowds. Fourth, our study focused on the cultural context of American misinformation, and recruited American participants. Future work should explore how such findings translate to other cultural contexts. In addition, our design cannot disentangle the effectiveness of warning labels attributed to the human-AI system, versus any warning label at all. We chose the human-AI system as a baseline because we are primarily focused on the role of explanations, but more work is necessary to distinguish the effects of more traditional warning labels (e.g. those from (Yaqub et al. 2020)) with the specialized labels from the human-AI system. Future work might do so by including controls for both a standard, non-attributed warning label, and the human-AI baseline we used. Our study also relies on perceptions of a *hypothetical* human-AI labeling system. It is unclear if participants believe that the warning labels actually came from a human-AI system, which may undermine our results. Correspondingly, we did not investigate how the act of sharing intersects the labeled content. For example, a user might be happy that a piece of content was labeled and might want to share the warning label itself with their friends, which in our data would be interpreted as the warning label not working. These quandaries underscore the challenges of understanding the nuances of genuine social media behavior in a survey context.

Taking a step back, it's important to note that our approach of design interventions to combat misinformation is but one tactic to fight a system with many actors and perverse incentives. Our intervention may systematically benefit large companies by providing them "band-aid" solutions to the systemic problems of an attention economy, rather than forcing them to consider the underlying causes. Ultimately, we hope our work invigorates careful research on how conventional approaches to fighting misinformation and making systems transparent intersect with the subtle, complex and consequential psychological dynamics of social media.

References

- Abdali, S.; Gurav, R.; Menon, S.; Fonseca, D.; Entezari, N.; Shah, N.; and Papalexakis, E. E. 2021. Identifying Misinformation from Website Screenshots. *arXiv preprint arXiv:2102.07849*.
- Allen, J.; Martel, C.; and Rand, D. 2021. Birds of a feather

- don't fact-check each other: Partisanship and the evaluation of news in Twitter's Birdwatch crowdsourcing fact-checking program. In *Forthcoming in Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–11.
- Allen, J. N. L.; Arechar, A. A.; Pennycook, G.; and Rand, D. 2020. Scaling up fact-checking using the wisdom of crowds. Bagherpour, A. 2020. COVID Misinformation Is Killing People. <https://www.scientificamerican.com/article/covid-misinformation-is-killing-people1/>, Accessed: 2021-12-12.
- Bansal, G.; Wu, T.; Zhu, J.; Fok, R.; Nushi, B.; Kamar, E.; Ribeiro, M. T.; and Weld, D. S. 2020. Does the whole exceed its parts? the effect of ai explanations on complementary team performance. *arXiv preprint arXiv:2006.14779*.
- Berinsky, A. J.; Margolis, M. F.; Sances, M. W.; and Warshaw, C. 2021. Using screeners to measure respondent attention on self-administered surveys: Which items and how many? *Political Science Research and Methods*, 9(2): 430–437.
- Bode, L.; and Vraga, E. K. 2015. In related news, that was wrong: The correction of misinformation through related stories functionality in social media. *Journal of Communication*, 65(4): 619–638.
- Bozarth, L.; and Budak, C. 2020. Toward a better performance evaluation framework for fake news classification. In *Proceedings of the international AAAI conference on web and social media*, volume 14, 60–71.
- Bozarth, L.; Saraf, A.; and Budak, C. 2020. Higher ground? How groundtruth labeling impacts our understanding of fake news about the 2016 US presidential nominees. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, 48–59.
- Breuer, A.; Eilat, R.; and Weinsberg. 2020. Detecting fake accounts on social networks with SybilEdge. <https://research.fb.com/blog/2020/04/detecting-fake-accounts-on-social-networks-with-sybiledge/>, Facebook Research Blog.
- Brooks, S. P.; and Gelman, A. 1998. General methods for monitoring convergence of iterative simulations. *Journal of computational and graphical statistics*, 7(4): 434–455.
- Bürkner, P.-C. 2017. brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software, Articles*, 80(1): 1–28.
- Carpenter, B.; Gelman, A.; Hoffman, M. D.; Lee, D.; Goodrich, B.; Betancourt, M.; Brubaker, M.; Guo, J.; Li, P.; and Riddell, A. 2017. Stan: A Probabilistic Programming Language. *Journal of Statistical Software, Articles*, 76(1): 1–32.
- Chan, M.-p. S.; Jones, C. R.; Hall Jamieson, K.; and Albarracín, D. 2017. Debunking: A meta-analysis of the psychological efficacy of messages countering misinformation. *Psychological science*, 28(11): 1531–1546.
- Clayton, K.; Blair, S.; Busam, J. A.; Forstner, S.; Gance, J.; Green, G.; Kawata, A.; Kovvuri, A.; Martin, J.; Morgan, E.; et al. 2019. Real solutions for fake news? Measuring the effectiveness of general warnings and fact-check tags in reducing belief in false stories on social media. *Political Behavior*, 1–23.
- Coleman, K. 2021. Introducing Birdwatch, a community-based approach to misinformation. https://blog.twitter.com/en_us/topics/product/2021/introducing-birdwatch-a-community-based-approach-to-misinformation.html, Twitter Blog.
- Constance, J. 2016. Facebook Now Flags and Down-Ranks Fake News with Help from Outside Fact Checkers. *TechCrunch, TechCrunch*, 15.
- Coppock, A.; and McClellan, O. A. 2019. Validating the demographic, political, psychological, and experimental results obtained from a new source of online survey respondents. *Research & Politics*, 6(1): 2053168018822174.
- Cramer, H.; Evers, V.; Ramlal, S.; Van Someren, M.; Rutledge, L.; Stash, N.; Aroyo, L.; and Wielinga, B. 2008. The effects of transparency on trust in and acceptance of a content-based art recommender. *User Modeling and User-adapted interaction*, 18(5): 455.
- Doshi-Velez, F.; and Kim, B. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.
- Epstein, Z.; Berinsky, A. J.; Cole, R.; Gully, A.; Pennycook, G.; and Rand, D. G. 2021. Developing an accuracy-prompt toolkit to reduce COVID-19 misinformation online. *Harvard Kennedy School Misinformation Review*.
- Epstein, Z.; Pennycook, G.; and Rand, D. 2020. Will the crowd game the algorithm? Using layperson judgments to combat misinformation on social media by downranking distrusted sources. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–11.
- Frederick, S. 2005. Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*, 19(4): 25–42.
- Garrett, R. K.; and Poulsen, S. 2019. Flagging Facebook falsehoods: Self-identified humor warnings outperform fact checker and peer warnings. *Journal of Computer-Mediated Communication*, 24(5): 240–258.
- Granik, M.; and Mesyura, V. 2017. Fake news detection using naive Bayes classifier. In *2017 IEEE First Ukraine Conference on Electrical and Computer Engineering (UKRCON)*, 900–903. IEEE.
- Horne, B. D.; Nevo, D.; O'Donovan, J.; Cho, J.-H.; and Adalı, S. 2019. Rating reliability and bias in news articles: Does ai assistance help everyone? In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 13, 247–256.
- Kane, J. V.; and Barabas, J. 2019. No harm in checking: Using factual manipulation checks to assess attentiveness in experiments. *American Journal of Political Science*, 63(1): 234–249.
- Kim, J.; Tabibian, B.; Oh, A.; Schölkopf, B.; and Gomez-Rodriguez, M. 2018. Leveraging the crowd to detect and reduce the spread of fake news and misinformation. In *Proceedings of the eleventh ACM international conference on web search and data mining*, 324–332.
- Kizilcec, R. F. 2016. How much information? Effects of transparency on trust in an algorithmic interface. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 2390–2395.

- Kulesza, T.; Stumpf, S.; Burnett, M.; Yang, S.; Kwan, I.; and Wong, W.-K. 2013. Too much, too little, or just right? Ways explanations impact end users' mental models. In *2013 IEEE Symposium on Visual Languages and Human Centric Computing*, 3–10. IEEE.
- Lai, V.; and Tan, C. 2019. On human predictions with explanations and predictions of machine learning models: A case study on deception detection. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 29–38.
- Langer, E. J.; Blank, A.; and Chanowitz, B. 1978. The mindlessness of ostensibly thoughtful action: The role of "placebic" information in interpersonal interaction. *Journal of personality and social psychology*, 36(6): 635.
- Lewandowsky, S.; Cook, J.; Ecker, U.; Lewandowsky, S.; Cook, J.; Ecker, U.; and Newman, E. 2021. Under the Hood of The Debunking Handbook 2020: A consensus-based handbook of recommendations for correcting or preventing misinformation.
- Lewandowsky, S.; Ecker, U. K.; Seifert, C. M.; Schwarz, N.; and Cook, J. 2012. Misinformation and its correction: Continued influence and successful debiasing. *Psychological science in the public interest*, 13(3): 106–131.
- Liao, Q. V.; Gruen, D.; and Miller, S. 2020. Questioning the AI: informing design practices for explainable AI user experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–15.
- Lipton, Z. C. 2018. The mythos of model interpretability. *Queue*, 16(3): 31–57.
- Lyons, T. 2021. Increasing Our Efforts to Fight False News. <https://about.fb.com/news/2018/06/increasing-our-efforts-to-fight-false-news/>, Facebook Newsroom.
- Martel, C.; Mosleh, M.; and Rand, D. G. 2021. You're definitely wrong, maybe: Correction style has minimal effect on corrections of misinformation online. *Media and Communication*.
- Mayer, R. C.; Davis, J. H.; and Schoorman, F. D. 1995. An integrative model of organizational trust. *Academy of Management Review*, 20: 709–734.
- Morrow, G.; Swire-Thompson, B.; Polny, J.; Kopec, M.; and Wihbey, J. 2020. The Emerging Science of Content Labeling: Contextualizing Social Media Content Moderation. Available at SSRN.
- Mosleh, M.; Pennycook, G.; and Rand, D. G. 2020. Self-reported willingness to share political news articles in online surveys correlates with actual sharing on Twitter. *Plos one*, 15(2): e0228882.
- Mosseri, A. 2017. Working to stop misinformation and false news. *Facebook Newsroom*, 7.
- Papenmeier, A.; Englebienne, G.; and Seifert, C. 2019. How model accuracy and explanation fidelity influence user trust. *arXiv preprint arXiv:1907.12652*.
- Pennycook, G.; Bear, A.; Collins, E. T.; and Rand, D. G. 2020. The implied truth effect: Attaching warnings to a subset of fake news headlines increases perceived accuracy of headlines without warnings. *Management Science*.
- Pennycook, G.; Binnendyk, J.; Newton, C.; and Rand, D. G. 2021a. A practical guide to doing behavioral research on fake news and misinformation. *Collabra: Psychology*, 7(1): 25293.
- Pennycook, G.; Cannon, T. D.; and Rand, D. G. 2018. Prior exposure increases perceived accuracy of fake news. *Journal of experimental psychology: general*, 147(12): 1865.
- Pennycook, G.; Epstein, Z.; Mosleh, M.; Arechar, A. A.; Eckles, D.; and Rand, D. G. 2021b. Shifting attention to accuracy can reduce misinformation online. *Nature*, 1–6.
- Pennycook, G.; and Rand, D. 2021a. Reducing the spread of fake news by shifting attention to accuracy: Meta-analytic evidence of replicability and generalizability. <https://psyarxiv.com/v8ruj/>. PsyArXiv.
- Pennycook, G.; and Rand, D. G. 2019. Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences*, 116(7): 2521–2526.
- Pennycook, G.; and Rand, D. G. 2021b. The psychology of fake news. *Trends in cognitive sciences*.
- Rader, E.; Cotter, K.; and Cho, J. 2018. Explanations as mechanisms for supporting algorithmic transparency. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, 1–13.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. "Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144.
- Roose, K. 2018. We Asked for Examples of Election Misinformation. You Delivered. *nytimes.com*.
- Rosenfeld, N.; Szanto, A.; and Parkes, D. C. 2020. A kernel of truth: Determining rumor veracity on twitter by diffusion pattern alone. In *Proceedings of The Web Conference 2020*, 1018–1028.
- Roth, Y.; and Pickles, N. 2020. Updating our approach to misleading information. *Twitter Blog*.
- Rubin, V. L.; Chen, Y.; and Conroy, N. K. 2015. Deception detection for news: three types of fakes. *Proceedings of the Association for Information Science and Technology*, 52(1): 1–4.
- Seo, H.; Xiong, A.; and Lee, D. 2019. Trust It or Not: Effects of Machine-Learning Warnings in Helping Individuals Mitigate Misinformation. In *Proceedings of the 10th ACM Conference on Web Science*, 265–274.
- Sherman, I. N.; Redmiles, E. M.; and Stokes, J. W. 2020. Designing Indicators to Combat Fake Media. *arXiv preprint arXiv:2010.00544*.
- Silverman, H. 2019. Helping Fact-Checkers Identify False Claims Faster. <https://about.fb.com/news/2019/12/helping-fact-checkers/>. Accessed: 2021-2-22. Facebook Research Blog.
- Thomson, K. S.; and Oppenheimer, D. M. 2016. Investigating an alternate form of the cognitive reflection test. *Judgment and Decision making*, 11(1): 99.

- Vosoughi, S.; Roy, D.; and Aral, S. 2018. The spread of true and false news online. *Science*, 359(6380): 1146–1151.
- Vraga, E. K.; and Bode, L. 2017. Using expert sources to correct health misinformation in social media. *Science Communication*, 39(5): 621–645.
- Walker, M.; and Gottfried, J. 2019. Republicans far more likely than Democrats to say fact-checkers tend to favor one side. *Pew Research Center*.
- Weinzierl, M.; Hopfer, S.; and Harabagiu, S. M. 2021. Misinformation Adoption or Rejection in the Era of COVID-19. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, 787–795.
- Yaqub, W.; Kakhidze, O.; Brockman, M. L.; Memon, N.; and Patil, S. 2020. Effects of Credibility Indicators on Social Media News Sharing Intent. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–14.
- Zannettou, S. 2021. "I Won the Election!": An Empirical Analysis of Soft Moderation Interventions on Twitter. *arXiv preprint arXiv:2101.07183*.
- Zhang, Y.; Liao, Q. V.; and Bellamy, R. K. 2020. Effect of confidence and explanation on accuracy and trust calibration in ai-assisted decision making. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 295–305.