Repetition Increases both the Perceived Truth and Fakeness of Information: An Ecological Account

Article *in* Cognition · September 2020

CITATIONS
0

READS
632

3 authors:

Olivier Corneille
Université Catholique de Louvain - UCLouvain
137 PUBLICATIONS **4,005** CITATIONS
SEE PROFILE

Adrien Mierop
Université Catholique de Louvain - UCLouvain
18 PUBLICATIONS **129** CITATIONS
SEE PROFILE

Christian Unkelbach
University of Cologne
98 PUBLICATIONS **2,693** CITATIONS
SEE PROFILE

Repetition Increases both the Perceived Truth and Fakeness of Information:

An Ecological Account

Olivier Corneille, Adrien Mierop

UCLouvain

&

Christian Unkelbach

University of Cologne

Manuscript in press at *Cognition*

Abstract

People believe repeated statements more compared to new statements – they show a truth by repetition effect. In three pre-registered experiments, we show that repetition may also increase perceptions that statements are used as fake news on social media, irrespective of the factual truth or falsehood of the statements (Experiment 1 & 2), but that repetition reduces perceptions of falsehood when the context of judgment is left unspecified (Experiment 3). On a theoretical level, the findings support an ecological account of repetition effects, as opposed to either a fluency-as-positivity or to an amplification account of these effects. On a practical level, they qualify the influence of repetition on the perception of fake news.

Repetition Increases both the Perceived Truth and Fakeness of Information: An Ecological Account.

People judge repeated information truer as compared to new information. This effect is called the "illusory truth effect", "repetition-induced truth effect", or "truth-by-repetition effect" (see Unkelbach et al., 2019). Since its initial demonstration by Hasher and colleagues (1977), this effect has been frequently replicated (see Dechêne et al., 2010, for a quantitative review). One theoretical explanation for this effect is that repetition increases processing fluency (e.g., Reber & Schwarz, 1999; Unkelbach, 2007). In turn, because processing fluency and subjective truth are linked, repetition increases the subjective truth of statements (Arkes et al., 1989; Fazio et al., 2015). However, the relation between processing fluency and subjective truth is not fixed. The interpretation of fluency experiences may be shaped by learning and context (e.g., Brinol, Petty, & Tormala, 2006; Unkelbach, 2006), and it is sensitive to people's naive theories (e.g., Winkielman & Schwarz, 2001) and causal models (e.g., Oppenheimer, 2004).

Based on Brunswik's (1955, 1957) lens model, Unkelbach and Greifeneder (2013) proposed an ecological model that conceptualizes fluency as a cue that people use across a broad range of judgments. According to this model, the interpretation of fluency cues is not fixed. Instead, people interpret processing fluency in accordance with the judgment context and their learning history; for example, within the context of a murder mystery novel, the easy and fluent solution is typically wrong. This model explains the influence of processing fluency on multiple judgment dimensions (e.g., Mandler et al., 1987), and it also explains reversals of fluency effects (e.g., Brinol et al., 2006; Oppenheimer, 2004; Unkelbach, 2006).

Predicting a *fake news by repetition* effect:

One seemingly paradoxical implication of this model is that repetition-induced fluency might, depending on the judgment context, increase feelings of truth or feelings of fakeness.

Unkelbach and Greifeneder (2013)'s model posits that there is nothing that feels inherently true in repeated statements. Because truth can be reasonably assumed to be more prevalent than falsehood in the typical information ecology (see Unkelbach, Koch, & Alves, 2019), people may use fluency as a cue for truth when the judgment context is unspecified (i.e., a truth-by-repetition effect; see Reber & Unkelbach, 2010). In other information ecologies[1], though, people may just use the cue differently. In particular, because people now have been familiarized with the widespread diffusion of fake news on social media, they may use their fluency feeling as a cue that a piece of information is actually "fake news" when the context of judgment relates to social media.

Consistent with this reasoning, Vosoughi, Roy, and Aral (2018) showed for a large sample of fact-checked information on Twitter that fake news in the scientific, urban legends, and political domains spread deeper (i.e., how often is information repeated/retweeted by unique users) and broader (i.e., how many users are involved over time) on social media. In addition, fake news' maximum depth and breadth is achieved faster. As a result, people may consider fluency due to repetition as a cue for fake news rather than a cue for truth in the information ecology of social media.

Using fluency as a cue for fake news does not imply that people's accuracy in making these judgments will necessarily increase. In Brunswik's lens-model terms, this would be a question of the cue's validity (Kareleia & Hogarth, 2008). Yet, the model predicts variations in the interpretation of the fluency cue across judgment contexts (see Unkelbach & Greifeneder, 2013, p. 14); we are addressing the cue's interpretation, not the cue's validity. We anticipated that, in the context of fake news on social media, people may interpret fluency as a cue that a given information might be fake news. If this is the case, repeated information may be more readily considered "fake news" in a social media context, constituting a *fake news by repetition* effect.

Alternative Explanations and differential predictions

The model by Unkelbach and Greifeneder (2013), which serves as the basis for the main prediction, may be contrasted to two alternative accounts of the truth-by-repetition effect that make different predictions. Table 1 summarizes these divergent predictions.

Table 1. Differential predictions for the effect of repetition on judgments of truth, falsehood, and use as fake news on social media, from the perspective of a Fluency-Positivity account, an Amplification account, and an Ecological account.

	Judgments		
	"True"	"False"	"Fake News"
Context	(Unspecified)	(Unspecified)	(Social Media)
Positivity Account	Increases	No effect/Reduces	No effect/Reduces
Amplification Account	Increases	Increases	Increases
Ecological Account	Increases	No effect/Reduces	Increases

Note. The "Truth" and "False" judgments refer to unspecified, and thus, general ecologies.

A first alternative account, which we may call the *positivity account*, posits that processing fluency is an inherently positive experience (e.g., Winkielman & Cacioppo, 2001; Winkielman et al., 2003); this results in assimilative judgments on positively-connotated dimensions such as "truth" (Unkelbach et al., 2011). Because "fake news" is negatively connotated, the positivity account would not predict a *fake news by repetition* effect. For the same reason, it would not predict repetition to increase judgments of falsehood in general. Possibly, this account would actually predict decreased judgments of fakeness and falsehood for repeated information. This is because what is considered truer/more positive should also be considered less false/less negative.

A second alternative account, the *amplification account*, has been proposed by Albrecht and Carbon (2014; see also Mandler et al., 1987). These authors argued that fluency might amplify judgments on any dimension (see also Landwehr & Eckmann, 2020; for a comparison of these types of fluency effects). This account would predict the same effect of repetition on any type of judgment, with repetition increasing judgments of truth, falsehood or fakeness (see Table 1, last row).

The present *ecological* account predicts that repetition increases judgments of truth (and possibly decreases judgment of falsehood) for unspecified judgment contexts, because truthful information should be prevalent across all information ecologies (Grice, 1975; Unkelbach, Koch, & Alves, 2019). In addition, it predicts a higher rate of classification of repeated statements as fake news in a social media context (see Table 1, first row), because in that specific ecology, fake news spread deeper, broader, and faster (see Vosoughi et al., 2018).

Although it may seem contradictory that fluency may result in a higher rate of classification of a statement as "true" (unspecified context) or "fake news" (social media context), the contradiction is more apparent than real and follows from an ecological cue model. Let us illustrate the ecology's importance with two further points. First, the same information, experiential or factual, changes its meaning as a function of the context, or more generally, the ecology. For example, people may find a sea fruit pizza appealing in general, but not so much in the context of an unsanitary restaurant. Second, the judgment also differs in our current question (i.e., "Is this true?" vs. "Is this fake news?"); while the question if something is true or not may prompt a general ecology (i.e., "Do you like sea fruit pizza?"), the question about fake news on social media prompts a very specific context (i.e., "Do you find sea fruit pizza in this restaurant edible?")

We reasoned that people's use of the fluency cue on social media might be opposite to people's use of the cue in a general ecology. To stay with our example, they might like sea fruit pizza in general, but be very cautious about a sea fruit pizza at creepy restaurants. Likewise, repetition-induced fluency may increase perceptions that a statement is "true" when the ecology is left unspecified but may conversely increase perceptions that a statement is "fake news" when the ecology is that of social media.

Testing a repetition effect for judgments of truth, fake news, and falsehood provides at least two insights. Empirically, it is inherently interesting to test how repetition influences perceptions of fake news and falsehood, two dimensions that have not been investigated so far. Theoretically, as indicated by the unique predictions of the accounts across dimensions, the comparison of judgments across dimensions allows testing theoretical accounts of the truth-by-repetition effect.

The present research

Our primary goal was testing the *fake news by repetition* effect, and to contrast potential explanations for this anticipated effect. We report three pre-registered experiments that first test the *fake news by repetition* hypothesis (Experiment 1 & 2) and then a general amplification hypothesis (Experiment 3).

For Experiments 1 and 2, we changed the typical judgment context in studies on truth by repetition, from truth judgments to judgments for use as fake news. In addition, instead of the typical trivia statements (e.g., Hasher et al., 1977; Bacon, 1979), we used statements that participants may consider as relevant and interesting and might thus be seen as potential "fake news" (see below). We additionally examined whether this *fake news by repetition* effect, if found, would be moderated by the factual truth of the statements. We anticipated two possible outcomes. First, participants may possess some knowledge about the true statements, and they may use this knowledge to determine that factually true information cannot be "fake news".

This should lead to a larger effect of repetition for false compared to true statements. Given the existing literature of previous knowledge on subjective truth (Fazio et al., 2015), this seemed unlikely though. Alternatively, people may not have accessible knowledge about the true statements, but perhaps just weak subjective experiences with them. That is, these statements may be somewhat more fluently processed because they have been encountered before (see Unkelbach & Stahl, 2009, Exp. 2) or because they fit better with existing knowledge structures (Unkelbach & Rom, 2017). If this is the case, the expected *fake news by repetition* effect should not vanish for these statements, and so there should be no moderation of repetition effects by the factually true vs. false status of the statement. Experiment 3 then tests the possibility of a general amplification effect by asking for "false" judgments in unspecified contexts, rather than for "fake news" in a social media context. The pattern of results across experiments allows testing the presented competing theoretical explanations.

The pre-registration protocol, materials, experiment script, data, and analysis script for the first two experiments are available at this link: https://osf.io/v27mg/ and for the third experiment at this link: https://osf.io/5hsug/ These are the only three experiments we conducted on these questions so far. Experiment 1 shows the predicted *fake news by repetition* effect. Experiment 2 replicates this effect with a substantially increased sample size. Experiment 3 replicates the truth-by-repetition effect and also finds that repetition significantly *decreases* judgments of falsehood.

Experiment 1.

In Experiment 1, we used a typical truth-by-repetition paradigm (see Unkelbach, Koch, Silva, & Garcia-Marques, 2019). Participants read statements in a learning phase and evaluated these repeated and new statements in a later judgment phase. Yet, different from a typical truth-by-repetition experiment, we asked participants to judge whether they thought the statement had been used as fake news on the social media, for statements that participants

may see as relevant and interesting. The sampling of participants was unrestricted, except for the use of English language.

Participants and design

The truth-by-repetition effect is robust and may be found with less than 20 participants (e.g., Unkelbach, 2007). As we used a novel variation and only 40 statements (twenty factually true and twenty factually fake statements), we collected data from 49 participants (M_{age} = 35.74, SD_{age} = 12.44, 20 women) using the online Prolific Academic platform. We manipulated all factors within-participants: Repetition status (i.e., old vs. new) and factual truth (i.e., factually true vs. factually false).

Materials

We used 40 statements from Unkelbach and Speckman (2020). These statements were generated and pre-tested to be relevant and interesting. Half of these statements were factually true and half were factually false. The statements are available from the online materials. We conducted the experiment online, using the Prolific platform with a computer program written with OpenSesame (Mathôt et al., 2012).

Procedure

Participants first read a consent form informing them that the experiment was about judging statements. After reading the informed consent and agreeing to participate, the computer program started with the following instruction (adapted from Unkelbach & Rom, 2017): "Dear Participant, you are participating in a study investigating various influences on people's assessment of statements. Your task will be to judge whether statements have been previously used as a Fake News on the social media. In the first part, you will only see a presentation of statements. Please try to read all statements, even though the presentation is fast. Then, the judgment of statements will start. You should judge for each statement whether

it has been previously used as a Fake News on the social media by pressing a key. If you have no further questions, please commence with the presentation."

The presentation phase involved 20 statements (i.e., 10 true and 10 false) that were randomly selected for each participant from a set of 20 factually true statements (e.g., “The first heart transplantation took place in Vienna.”) and 20 factually false statements (e.g., “Russia produces the most aluminum worldwide.”). Statements were sequentially presented onscreen for 2500ms each, with a 1s break between statements.

The judgment phase started immediately after the presentation phase[2]. Participants judged 40 statements, 20 of them repeated, and 20 new; orthogonally, half were factually true or false. Statement presentation was randomized anew for each participant. The question “Do you believe this statement has been previously used as a Fake News on the social media” was visible on the upper half of the screen. Statements remained visible onscreen until participants responded (“Yes, the statement has been previously used as a Fake News on the social media”; “No, the statement has not been previously used as a Fake News on the social media”). Once a response was entered, the statement was removed from the screen and replaced by a blank screen for 1000ms, directly followed by a new trial. After completion of the study, participants were thanked for their participation and debriefed. They were invited to contact the person in charge of the research if they had any queries or comment on the study.

Results

We analyzed the data with a general linear model using the ‘lmer’ function from the ‘lme4’ packages from R (Bates et al., 2007). We used the proportion of "fake news" judgments (similar to PTJs; see Unkelbach & Rom, 2017; Unkelbach & Greifeneder, 2018, Experiment 1) as our main dependent variable. We analyzed this proportion as a function of repetition (i.e., old vs. new) and factual truth (i.e., factually true vs. false), and their interaction (see Figure 1’s left panel for means and distributions).

As predicted, participants judged old statements more often as fake news (M = .52; SD = .30) than new statements (M = .42; SD = .23), $F(1, 143.06) = 10.06$, $p = .002$, $\eta^2_p = .059$. There was no significant main effect of factual truth. Participants descriptively judged factually true statements (M = .49; SD = .25) more frequently as "fake news" than factually false statements (M = .45; SD = .29), but this difference was not significant, $F(1, 143.06) = 3.52$, $p = .063$, $\eta^2_p = .017$. The interaction was not significant either, $F(1,190) = 1.79$, $p = .131$, $\eta^2_p = .009$.

Discussion

Experiment 1 establishes a *fake news by repetition* effect. In addition, this effect was not moderated by the factual truth of the statements. Because of the ambiguous level of significance of the main effect of truth, and to check for the robustness of the *fake news by repetition* effect, we replicated Experiment 1 with a substantially larger sample.

Experiment 2

Experiment 2 was a high-powered robustness test of Experiment 1.

Participants, Design, Materials, and Procedure

Experiment 2 replicated Experiment 1 exactly, but this time tripled the participants and collected data from 152 participants ($M_{age} = 34.20$, $SD_{age} = 11.42$, 77 women).[3]

Results

We fully replicated Experiment 1, with a more clear-cut pattern of findings (see Figure 1's right panel). We replicated the *fake news by repetition* effect: participants judged old statements more often as "fake news" (M = .56; SD = .31) than new statements (M = .41; SD = .24), $F(1,448) = 22.13$, $p < .001$, $\eta^2_p = .045$. There was no reliable difference between factually true (M = .50; SD = .26) and factually false (M = .47; SD = .31) statements, $F(1,444.80) = 0.24$, $p = .62$, $\eta^2_p = .000$. There was also no interaction, $F(1, 446.42) = .22$, $p = .64$, $\eta^2_p = .000$.

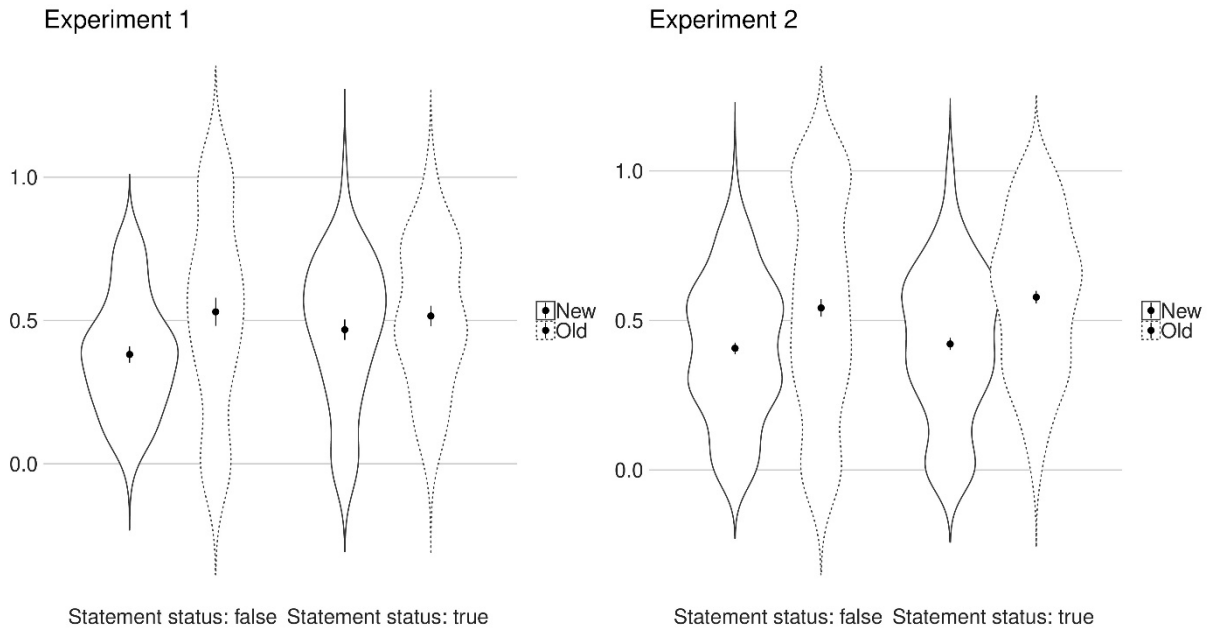

Figure 1. Proportion of "fake news" judgments as a function of Repetition status (Old vs. New) and Truth status (False vs. True). The violin plots represent the kernel probability density of the data at different values. Hence, they may exceed the minimum and maximum values of the individual data points (here, respectively 0 and 1).

Experiment 3

At this point, we have established a highly interesting empirical pattern: people judge repeated information more frequently as "fake news" compared to new information. Besides the practical implication, this finding is also inconsistent with a fluency-positivity account of the truth-by-repetition effect (see Table 1). However, as discussed, the *fake news by repetition* effect found in Experiments 1 and 2 is consistent with an amplification account (Albrecht & Carbon, 2014); fluency may increase the probability of agreeing with any kind of judgment at hand. To test this alternative account, Experiment 3 used a Repetition status (Old vs. New) by Question format ("Is this statement a true statement?" vs. "Is this statement a false statement?) mixed design with "yes" and "no" as the possible answers. We departed from the "Is this statement False or True?" format typically used in truth by repetition studies. This change follows from our research questions. Participants should frame the question *either* in

terms of truth *or* falsehood. In addition, a "False or True?" format would be confusing in the falsehood format condition, as participants may select the "No" response option either to indicate that the statement is false or that it should *not* be considered false (i.e., a double negation). If fluency amplifies confirmatory responses on any dimension, we would expect to replicate the truth-by-repetition effect and also find increases in judged falsehood for repeated statements. In contrast, if fluency is used as a context-specific cue, we would expect to replicate the truth-by-repetition effect, but we would *not* expect increases in judged falsehood after repeated exposure to a statement.

Participants, Design, Materials, and Procedure

Experiment 3's procedure and materials were similar to Experiment 1 and 2, but now opposed "True" vs. "False" response formats. As pre-registered, we collected data from 200 participants (M_{age} = 32.70, SD_{age} = 10.47, 100 women). Half of the participants had to judge if the statements were true whereas the other half of participant judged if the statements were false.

In the Truth (Falsehood) format condition, participants received the following instructions "Dear Participant, you are participating in a study investigating various influences on people's assessment of statements. Your task will be to judge whether statements are true (false). In the first part, you will only see a presentation of statements. Please try to read all statements, even though the presentation is fast. Then, the judgment of statements will start. You should judge whether statements are true (false) statements by pressing a corresponding key."

After the presentation of the statement and before the judgment of the statements, participants were reminded of the instructions: "The judgment of statements will now start. You should judge, for each statement, whether this is a true (false) statement by pressing a

corresponding key. To indicate that, YES, the statement is a true (false) statement, press "y". To indicate that, NO, the statement is not a true (false) statement, press "n".”

Results

To reiterate, the amplification account predicts a main effect of repetition independent of the judgment dimension. The ecological model predicts that without further context specification, repetition should increase judgments of truth, but decrease judgments of falsehood, leading to an interaction. Figure 2 presents participants’ average PTJ as a function of repetition and question format. As preregistered, we analyzed these data with a Repetition status (Old vs. New) by Question format (Truth vs. Falsehood) mixed ANOVA with repeated measures on the first factor.

As Figure 2 suggests, the interaction between Repetition status (old vs. new) and Question format (Truth vs. Falsehood) was highly significant, $F(1, 198) = 37.37$, $p < .001$, $\eta^2_p = .155$. Within the “truth” condition, participants judged old statements more often as true (M = .69; SD = .18) than new statements (M = .51; SD = .17), $t(104) = - 6.87$, $p < .001$, $\eta^2_p = .306$. This pattern reversed in the “falsehood” condition, where participants judged old statements *less* often as false (M = .40; SD = .24) than new statements (M = .48; SD = .16), $t(94) = 2.52$, $p = .013$, $\eta^2_p = .053$. Hence, repetition increased judgments of truth for the 'true' format but decreased judgments of falsehood for the 'false' format. Interestingly, these effects of repetition were significantly different in absolute value, $F(1, 198) = 5.91$, $p = .016$, $\eta^2_p = .024$, with a larger effect of repetition for the 'true' than for the 'false' format. As previously, factual truth of the statements did not influence participants’ responses. (all Fs < 1).

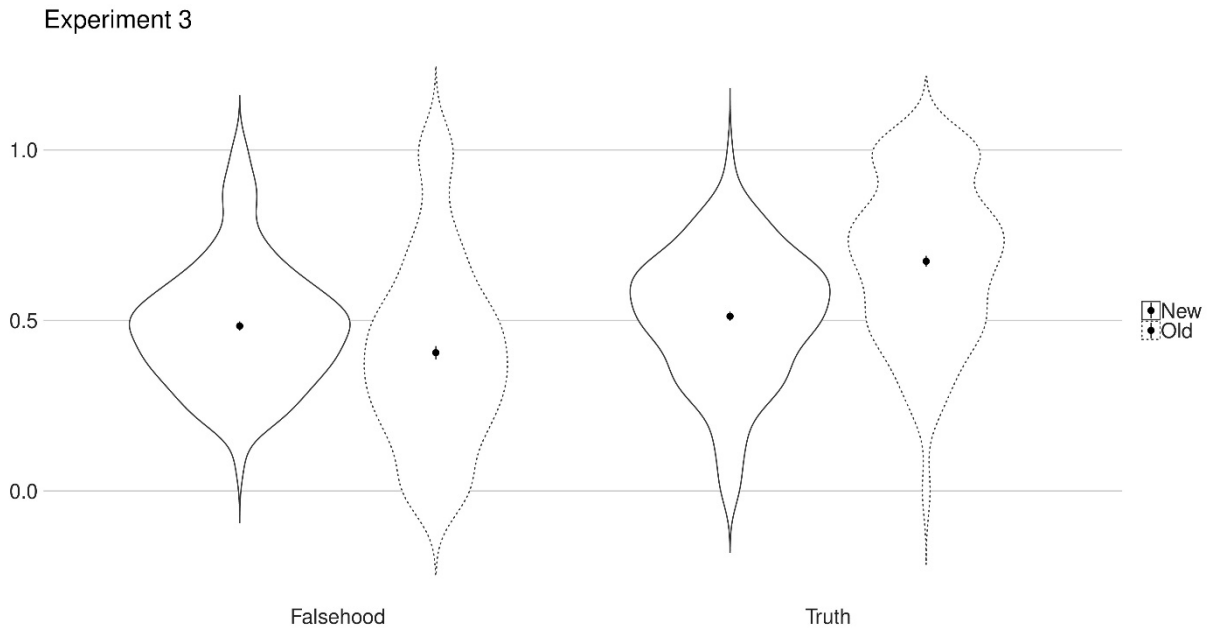

Figure 2. Proportion of "yes" judgments as a function of Repetition status (old vs. new) and Response format (Truth vs. Falsehood). The violin plots represent the kernel probability density of the data at different values.

General discussion

Judging whether information is true or false is a central task in a world where information is more accessible than ever before. A typical finding in such judgments of truth is that people believe repeated information more compared to new information (Hasher et al., 1977), and processing fluency is a prime theoretical construct that may underlie this repetition effect (Dechêne et al., 2010). Based on a general model of fluency effects (Unkelbach & Greifeneder, 2013), we hypothesized that in the information ecology of social media, repetition-induced fluency may be interpreted as a cue for "fake news".

We presented three pre-registered experiments that tested this hypothesis of a *fake news by repetition* effect and compared it with two other accounts: the hedonic fluency model (Winkielman & Cacioppo, 2001) and the amplification account (Albrecht & Carbon, 2014). The first two experiments established that participants categorize repeated statements more likely as "fake news" than new statements, supporting a *fake news by repetition* effect. The

result that repetition increases feelings of fakeness is inconsistent with a positivity account of the truth-by-repetition effect but is consistent with either an ecological or an amplification account of this effect. The third experiment, however, ruled out the amplification account. Here, repetition increased perceptions of truth but decreased perceptions of falsehood. As discussed in the introduction, this is consistent with either a positivity or an ecological account of repetition effects, but inconsistent with an amplification account of repetition effects. Across experiments, the pre-registered pattern of findings observed here is only consistent with an ecological account of repetition effects (see Table 1).

Nevertheless, the ecological account is not incompatible with a positivity account, as a cue might be inherently neutral or inherently positive. One might predict, however, that if the cue and criterion are evaluatively similar, the cue should receive more weight, a pattern we observed in Experiment 3. In addition, the data is also in line with people's use of naïve theories (e.g., Winkielman & Schwarz, 2001), as we did not test the learning of a cue, but only provided specific contexts (see also below).

In line with previous research, these effects were not qualified by the factual truth or falsehood of the statements. This absence of an interaction effect with statement's factual truth is unsurprising, in hindsight. Prior research found little (Unkelbach & Stahl, 2009) to no effects (Fazio et al., 2015; Fazio, Rand, & Pennycook, 2019) of existing knowledge. In addition, participants had little knowledge about the statements' factual truth to begin with. Likewise, the lack of a main effect for factual truth, that is, a fluency effect due to "natural", not experimental exposure, follows if these statements are unknown to most participants.

Interestingly, we found in Experiment 3 that repetition had a larger effect on perceptions of truth than on perceptions of falsehood. Considering the straightforward response format used in that experiment, this indicates that these judgments are not perfectly antagonistic. From an ecological perspective, it makes sense to posit that fluency is more

diagnostic of truth than falsehood: as truth is prevalent in the word, it may allow for a larger sampling basis and more possibilities for learning. This asymmetry may also be explained by the independent contribution of an amplification effect. Because responses are determined by multiple processes, repetition may induce both higher feelings of truth and higher amplification tendencies. In the truth format condition, these independent effects of repetition may go in the same direction, whereas in the falsehood format condition they may oppose each other, resulting in milder repetition effect. Finally, this differential effect might also indicate that the fluency experience is positively connotated, which aligns and magnifies judgments on the positive truth dimension but misaligns and attenuates judgments on the negative falsehood dimension (see above).

Importantly, one should not conclude from the current findings that repetition will always and univocally increase perception of information as "fake news". Such overgeneralization would be perfectly opposed to our claim that fluency effects depend on the judgment context. We specifically hypothesized and found that participants use fluency as a cue for "fake news" on social media, which was our relevant ecological context. In addition, it would be interesting to compare the magnitude of the *fake news by repetition* effect in various population samples. For instance, younger and older social media users may differ in their social media experiences, and this may moderate the current effects.

The *fake news by repetition* effect established here has both theoretical and practical implications. Theoretically, this effect supports the view that meta-cognitive experiences such as fluency resulting from repetition may inform not just various, but sometimes also near opposite judgments. This is in line with a variety of findings showing that meta-cognitive experiences are malleable and depend on the context of judgment (e.g., Winkielman & Schwarz, 2001; Briñol et al., 2006; Hertwig et al., 2008). This indicates the adaptive and

flexible nature of the inferences people draw from their meta-cognitive experiences (e.g., Bjork et al., 2013; Koriat & Levy-Sadot, 1999).

Practically, the *fake news by repetition* effect qualifies the influence of fake news repetition. Past research indicates that people have a difficult time correcting for the influence of fake news that are later proved incorrect (e.g., Johnson, & Seifert, 1994, Lewandowsky et al., 2012). Likewise, explicit warnings do not eliminate the continued influence of misinformation (e.g., Ecker et al., 2010). Recent research also shows that repetition increases the perceived accuracy of fake news (Pennycook et al., 2018). This suggests that the repetition of fake news generates feelings of truth that are difficult to prevent and to correct. Whereas these conclusions can be held with confidence, the current findings indicate that repetition – at least under the minimal study conditions involved here – may also have the potential to increase feelings of untrustworthiness. Therefore, one strategy for tempering the negative impact of misinformation may be to orient individuals towards evaluations of fakeness instead of evaluations of truth.

Finally, an important epistemic question from an ecological perspective is whether the use of fluency positively or negatively influences judgment accuracy. In order to answer this question, one would need to know the "objective" state of the world. In the present case, one would need to ascertain that truth is prevalent in general, which is strongly suggested by the communication rules suggested by Grice (1975), whilst establishing that fake news is prevalent on social media. As done in the study by Voroughi et al. (2018), this would require large samples of available information and establish its factual truth status. At present, we are not aware of such a data set.

The theoretical rationale we proposed here, however, requires no such strong prevalence assumption. According to Unkelbach and Greifeneder (2013)'s general model of fluency, it suffices that participants learn a fluency/truth relation in general, but a

fluency/“fake news” relation in social media contexts. In addition, people may acquire naïve beliefs about the prevalence of fake news on the media, which would also lead to varying interpretations of fluency (e.g., Winkielman & Schwarz, 2001; Briñol et al., 2006). Future research may clarify whether such naïve theories are consciously held and are verbally reportable and, if they are, if they moderate the magnitude of *truth by repetition* and *fake news by repetition* effects. Such evidence may be seen as supportive of the use of deliberate inferences in these effects. Independent of the epistemological justification, the current study shows that fake news has become sufficiently widespread on the social media that the participants’ sample recruited used repetition and the resulting fluency as a cue to judge information as “fake news”.

Conclusion

A large body of literature supports the view that repetition increases perception of truth. We established that repetition may as well increase perceptions of fakeness (Experiments 1 and 2), even for factually true statements: a *fake news by repetition* effect. This finding cannot be reduced to the fact that fluency amplifies any kind of judgmental tendencies (Experiment 3). Rather, it is in line with an ecological model in which people use fluency as a cue whose interpretation varies, occasionally to antagonistic poles, as a function of the judgment context.

Footnotes

1. Consistent with Unkelbach and Greifeneder (2013), “ecologies” refer in this article to the contexts of judgment: “The fluency influence depends on the ecology in which a judgment is made; in experimental tasks, the ecology is often provided by the questions researchers are asking." (Unkelbach and Greifeneder, 2013; p. 20-21)

2. The pre-registration states that the judgment phase follows “a short distraction task”. In order to cut down experimental costs, we decided to remove the distraction task but we omitted to correct the pre-registration accordingly.

3. We originally planned to collect data on 170 participants but 18 participants from the Prolific Academic platform used their phone or their tablet rather than their computer as it was specified. Their data were therefore not encoded.

References:

Albrecht, S., & Carbon, C. C. (2014). The Fluency Amplification Model: Fluent stimuli show more intense but not evidently more positive evaluations. *Acta Psychologica*, *148*, 195-203.

Arkes, H. R., Hackett, C., & Boehm, L. (1989). The generality of the relation between familiarity and judged validity. *Journal of Behavioral Decision Making, 2*, 81–94.

Bacon, F. T. (1979). Credibility of repeated statements: Memory for trivia. *Journal of Experimental Psychology: Human Learning and Memory, 5*, 241–252.

Bates, D., Sarkar, D., Bates, M. D., & Matrix, L. (2007). The lme4 package. R package version, 2(1), 74.

Bjork, R. A., Dunlosky, J., & Kornell, N. (2013). Self-regulated learning: Beliefs, techniques, and illusions. *Annual Review of Psychology, 64*, 417-444.

Briñol, P., Petty, R. E., & Tormala, Z. L. (2006). The malleable meaning of subjective ease. *Psychological Science, 17*, 200-206.

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review , 62*, 193-217.

Brunswik, E. (1957). Scope and aspects of the cognitive problem . In H. Gruber , K. R. Hammond, & R. Jessor (Eds.), *Contemporary approaches to cognition* (pp. 5 – 31). Cambridge, MA: Harvard University Press.

Dechêne, A., Stahl, C., Hansen, J., & Wänke, M. (2010). The truth about the truth: A meta-analytic review of the truth effect. *Personality and Social Psychology Review*, *14*, 238–257.

Ecker, U. K., Lewandowsky, S., & Tang, D. T. (2010). Explicit warnings reduce but do not eliminate the continued influence of misinformation. *Memory & Cognition*, *38*, 1087-1100. doi:10.3758/MC.38.8.1087

Johnson, H. M., & Seifert, C. M. (1994). Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 1420–1436.

Fazio, L. K., Brashier, N. M., Payne, B. K., & Marsh, E. J. (2015). Knowledge does not protect against illusory truth. *Journal of Experimental Psychology: General, 144*, 993.

Fazio, L. K., Rand, D. G., & Pennycook, G. (2019). Repetition increases perceived truth equally for plausible and implausible statements. *Psychonomic Bulletin & Review*, *26*, 1705-1710.

Grice, H. P. (1975). Logic and conversation. In P. Cole and J. L. Morgan (Eds.) *Syntax and semantics: Speech acts* (Vol. 3; pp. 41-58). New York, NY: Seminar Press.

Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior, 16*, 107–112.

Hertwig, R., Herzog, S. M., Schooler, L. J., & Reimer, T. (2008). Fluency heuristic: A model of how the mind exploits a by-product of information retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*, 1191-1206.

Karelaia, N., & Hogarth, R. (2008). Determinants of linear judgment: A meta-analysis of lens model studies. *Psychological Bulletin, 134*, 404-426.

Koriat, A., & Levy-Sadot, R. (1999). Processes underlying metacognitive judgments: Information-based and experience-based monitoring of one's own knowledge. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (p. 483–502). New York, NY:Guilford Press.

Landwehr, J. R., & Eckmann, L. (2020). The nature of processing fluency: Amplification versus hedonic marking. *Journal of Experimental Social Psychology*, *90*, 103997.

Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, *13*, 106-131.

Mandler, G., Nakamura, Y., & Van Zandt, B. J. (1987). Nonspecific effects of exposure on stimuli that cannot be recognized. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 646–648.

Mathôt, S., Schreij, D., & Theeuwes, J. (2012). OpenSesame: An open-source, graphical experiment builder for the social sciences. *Behavior Research Methods, 44*(2), 314-324.

Morey, R. D., & Rouder, J. N. (2015). BayesFactor 0.9. 12-2. Comprehensive R Archive Network.

Oppenheimer, D. M. (2004). Spontaneous discounting of availability in frequency judgment tasks. *Psychological Science, 15*, 100–105.

Pennycook, G., Cannon, T. D., & Rand, D. G. (2018). Prior exposure increases perceived accuracy of fake news. *Journal of Experimental Psychology: General*, *147*(12), 1865.

Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition, 8*, 338–342.

Reber, R., & Unkelbach, C. (2010). The epistemic status of processing fluency as source for judgments of truth. *Review of Philosophy and Psychology, 1*, 563–581.

Unkelbach, C. (2006). The learned interpretation of cognitive fluency. *Psychological Science*, *17*(4), 339-345.

Unkelbach, C. (2007). Reversing the truth effect: Learning the interpretation of processing fluency in judgments of truth. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33*, 219–230.

Unkelbach, C., Bayer, M., Alves, H., Koch, A., & Stahl, C. (2011). Fluency and positivity as possible causes of the truth effect. *Consciousness and Cognition, 20*, 594–602.

Unkelbach, C., & Greifeneder, R. (2013). A general model of fluency effects in judgment and decision making. In C. Unkelbach & R. Greifeneder (Eds.), *The experience of thinking: How the fluency of mental processes influences cognition and behaviour* (pp. 11–32). New York, NY: Psychology Press.

Unkelbach, C., & Greifeneder, R. (2018). Experiential fluency and declarative advice jointly inform judgments of truth. *Journal of Experimental Social Psychology, 79*, 78-86.

Unkelbach, C., Koch, A., & Alves, H. (2019). The Evaluative Information Ecology: On the frequency and diversity of "good" and "bad". *European Review of Social Psychology, 30,* 216-270.

Unkelbach, C., Koch, A., Silva, R. R., & Garcia-Marques, T. (2019). Truth by repetition: Explanations and implications. *Current Directions in Psychological Science*, *28*(3), 247-253.

Unkelbach, C., & Rom, S. C. (2017). A referential theory of the repetition-induced truth effect. *Cognition, 160*, 110-126.

Unkelbach, C., & Speckmann, F. (2020). Costly beliefs in false information despite valid advice. Manuscript in preparation.

Unkelbach, C., & Stahl, C. (2009). A multinomial modeling approach to dissociate different components of the truth effect. *Consciousness and Cognition, 18*, 22–38.

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, *359*(6380), 1146-1151.

Winkielman, P., & Cacioppo, J. T. (2001). Mind at ease puts a smile on the face: psychophysiological evidence that processing facilitation elicits positive affect. *Journal of Personality and Social Psychology*, *81*(6), 989.

Winkielman, P., & Schwarz, N. (2001). How pleasant was your childhood? Beliefs about memory shape inferences from experienced difficulty of recall. *Psychological Science, 12*, 176-179.

Winkielman, P., Schwarz, N., Fazendeiro, T., & Reber, R. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. *The psychology of evaluation: Affective processes in cognition and emotion*, *189*, 217.