

Creeping Objectivity: Prior Exposure Leads People to Believe Claims Are More Objective

Daniel J. Mirny and Stephen A. Spiller

UCLA Anderson School of Management

Abstract

Whether people believe issues are matters of objective right and wrong or matters of subjective differences of opinion has important downstream consequences for judgment and interpersonal conflict. But perceived objectivity is a malleable construct, affected by how claims are presented. People can disagree about the same claims – one person will believe that a claim is objective (and can be fact-checked as either true or false) and another person will believe that the same claim is subjective (and not fact-checkable at all). Previous research has found that prior exposure increases the perceived veracity of objective claims (the illusory truth effect) as well as agreement with subjective claims (the mere exposure effect). The present research bridges these two literatures to investigate the novel question of whether prior exposure affects the perceived objectivity of claims. In a pre-registered experiment (N = 1,000 online participants), we find that prior exposure to claims results in people believing them to be more objective, across three different measures of perceived objectivity. We discuss potential processes for this creeping objectivity effect, along with theoretical and practical implications.

Keywords: prior exposure, objectivity, naïve realism, fact-check, opinions

All data, materials, pre-registrations, and supplementary materials are available at

https://researchbox.org/44&PEER_REVIEW_passcode=AAOJXL.

People encounter a variety of claims about the world in social media posts, conversations, news briefings, work emails, product reviews, and many other aspects of their daily lives. Some claims are factual statements (e.g., “It is raining”): they are objective and can be fact-checked as either true or false. Other claims are opinions (e.g., “The weather is nice”): they are subjective assessments, cannot be fact-checked, and other people may either agree or disagree with them. When a claim is perceived as an opinion, there is room for disagreement: Jack and Jill may both be aware that it is 95F in Boston, but Jack may believe that 95F is nice weather, Jill may believe that 95F is not nice weather, and even though they have divergent preferences, neither one of them is right nor wrong. However, when a claim is perceived as a factual statement, it necessitates the existence of only one, single, objectively correct view: Jack may believe that Bergen is the capital of Norway, Jill may believe that Oslo is the capital of Norway, and only one of them is right. As a result, perceived objectivity often stifles receptiveness to conflicting perspectives which are believed to be incorrect rather than reflecting a different opinion (Lieberman et al., 2012; Heiphetz & Young, 2017). The downstream consequences of perceived objectivity are severe, reaching into domains of our personal and professional lives, leading to interpersonal conflict (Ross & Ward, 1995), undermined collaboration (Lieberman et al., 2012), political partisanship (Skitka & Morgan, 2014; Blatz & Mercier, 2018), and moral tribalism (Johnson et al., 2021).

As a particularly stark example, perceived objectivity lies at the foundation of important policy and legal debates regarding the role of social and traditional media as credible information sources. For instance, in a defamation lawsuit brought against Tucker Carlson, Fox News’ attorney argued that Carlson was not liable due to the nature of his show being one of commentary, such that his statements “cannot reasonably be interpreted as facts” (*McDougal v. Fox News Network LLC*, 2020). A more recent case against former federal prosecutor Sidney Powell followed a similar precedent, with attorneys claiming that, “no reasonable person would conclude that the statements were truly statements of fact... Powell’s claims were her opinions” (*US Dominion, Inc. v. Powell*, 2021). The distinction between whether claims are classified as factual statements or as opinions also has substantive effects on the spread of

misinformation. Facebook's policy of not fact-checking opinion posts has stymied climate scientists' attempts to combat the spread of misinformation throughout the platform when climate change denialism posts are labeled as opinions (Penney, 2020).

These incidents emphasize the importance of understanding how people process and understand claim objectivity (Shane, 2017). What affects such assessments? Prior research has explored perceived objectivity as a feature of individual differences between people across cultures (Sarkissian et al., 2011; Goodwin & Darley, 2012) and across modes of argumentative interaction (Fisher et al., 2016). In the current research, we consider the effect of information presentation on perceived objectivity. Across a broad set of contexts, people frequently encounter the same claim multiple times. Bridging research on the illusory truth effect for objective claims (Hasher et al., 1977) and the mere exposure effect for subjective claims (Cacioppo & Petty, 1979), the current research investigates the effect of prior exposure to a claim on the perceived objectivity of that claim. In a three-stage experiment with a variety of different types of claims, participants are shown some claims repeatedly and other claims once. Using three different measures of perceived objectivity, we find that prior exposure to a claim increases how objective it is perceived to be. We also find that the effect of prior exposure on agreement with claims is correlated with the effect of prior exposure on the perceived objectivity of claims, but the effect of exposure on perceived objectivity does not represent a mere increase in agreement with claims.

There is Variability in the Perceived Objectivity of Claims

A recent Pew Research Center survey found that a majority of American respondents were unable to fully distinguish between factual statements and opinions in the news (Mitchell et al., 2018). Although respondents all saw the same set of claims, disagreement between respondents as to the objectivity of a given claim indicates that the perceived objectivity of a claim is not solely determined by the claim's content. This indication is supported by recent findings in linguistics, where even holding the content of a claim constant, a claim's linguistic structure can lead to different appraisals of whether the claim is more or less objective (Kaiser & Wang, 2020, 2021). But even holding the linguistic packaging constant,

beliefs about a claim's objectivity vary with one's other beliefs. People are prone to mischaracterizing the subjective nature of their own views and are more likely to believe that a claim is objective when it favors their own preferences (Griffin & Ross, 1991; Ross & Ward, 1995; Spiller & Belogolova, 2016; Toner et al., 2013) or reflects the perceived social and cultural consensus (Goodwin & Darley, 2008, 2012; Heiphetz & Young, 2017). This variability in perceived objectivity across people, contexts, and beliefs, suggests that perceived objectivity may also be malleable, influenced by the manner in which a claim is presented.

Effects of Prior Exposure on Factual Statements and Opinions

What might affect whether a claim is believed to be more or less objective? One potential but heretofore unexamined feature is whether or not the claim has been previously encountered. People often encounter the same claims repeatedly, and researchers have devoted considerable attention to the effect of prior exposure on beliefs about the veracity of facts and agreement with opinions. However, it may be the case that prior exposure can also affect higher-order beliefs, influencing how objective a claim is perceived to be.

Prior exposure increases the extent to which people believe factual statements to be true, deemed the illusory truth effect (Arkes et al., 1991; Begg et al., 1992; Dechêne et al., 2010; Hasher et al., 1977; Polage, 2012). This effect is robust and found across factual statements ranging from the mundane to the obscure (Bacon, 1970), from the self-evident to those contradicting existing knowledge (Fazio et al., 2015), and across consumer advertising, partisan political arguments, socio-political opinions, and fake news headlines (Arkes et al., 1989; Hawkins & Hock, 1992; Johar & Roggeveen, 2007; Pennycook et al., 2018). Moreover, Effron and Raj (2020) speculate that prior exposure may increase feelings of intuitive truthfulness, whereby information can feel true even when explicitly acknowledged as false (Shidlovski et al., 2014).

Prior exposure also increases the extent to which people agree with opinions and like different stimuli, deemed the mere exposure effect (Berryman, 1984; Bornstein, 1989; Crandall, 1985; Hill, 1978;

Pliner, 1982; Zajonc, 1968). Prior exposure affects the attitudes associated with opinions. When people have been previously exposed to an opinion, they are more likely to agree with it and attitudes associated with the opinion are more quickly recalled, demonstrate greater clarity, and are perceived to be more correct (Cacioppo & Petty, 1979; Petrocelli et al., 2007).

Both the illusory truth and mere exposure effects have been attributed to processing fluency, or ease of processing. Different operationalizations of fluency have resulted in similar effects, suggesting a common underlying cognitive mechanism for these and related effects (Alter & Oppenheimer, 2009; Begg et al, 1992; Jacoby et al. 1989; Oppenheimer, 2006; Reber et al, 1998; Shah & Oppenheimer, 2007; Thompson et al., 2013; Unkelbach, 2007; Unkelbach & Rom, 2017; Wang et al., 2016; Whittlesea, 1993).

Research on the effects of prior exposure has presented claims as either factual statements or as opinions, determined a priori, and has found that prior exposure increases (a) the belief that factual statements are accurate (illusory truth), and (b) agreement with opinions (mere exposure). Moreover, both effects can be at least partially explained by a shared mechanism of processing fluency. However, these results are focused on within-type malleability, that is, the malleable perceived accuracy of objective claims and the malleable agreement with subjective claims. As far as we are aware, research has not previously investigated whether the perceived objectivity of claims is *itself* affected by prior exposure.

We address this key unresolved question. While research on illusory truth and mere exposure has taken a claim's objectivity as given (objective in the illusory truth literature and subjective in the mere exposure literature), we bridge these findings to examine a cause of a claim's perceived objectivity. Research on naïve realism indicates that people sometimes treat their subjective assessments as though they are objective assessments (Griffin & Ross 1991; Ross & Ward 1995). Moral objectivity also varies across people, cultures, and modes of social interaction (Sarkissian et al., 2011; Goodwin & Darley, 2012; Fisher et al., 2016), suggesting that a claim's perceived objectivity is malleable, subject to the way in which it is presented. Given that (i) prior exposure is associated with increased agreement with claims (Hasher et al., 1977; Cacioppo and Petty, 1979), and (ii) strength of agreement is associated with

believing that claims are objective (Goodwin & Darley 2008; Spiller & Belogolova 2016; Toner et al., 2013), we hypothesize that prior exposure to a claim increases its perceived objectivity. We test this novel hypothesis in a three-stage experiment.

Method

In this experiment ($N = 1,000$), we examine the effect of prior exposure on the perceived objectivity of claims. This research was certified exempt by the home institution IRB. All data, materials (including all claims), and pre-registrations are available in a supplementary web appendix at https://researchbox.org/44&PEER_REVIEW_passcode=AAOJXL. In an effort to avoid a ‘file-drawer’ problem, the web appendix also includes a complete record (including data, materials, and analyses) of all preliminary pre-registered experiments in which we manipulated prior exposure and measured beliefs about claim objectivity.

Participants

We recruited a convenience sample of 1,000 participants (49.7% female, $M_{Age} = 40.4$) from Amazon Mechanical Turk using CloudResearch’s pool of “CloudResearch-Approved Participants” (Litman et al., 2020). Sample size was determined in advance and was large enough to provide greater than 85% power to detect a within-subject effect of at least 0.1 standard deviations in our target measure of perceived objectivity. This target effect size was based on preliminary studies reported in the supplemental materials.

Given the subtle manipulation, time delay, and an inability to control the experimental environment with remote participation, we were concerned about participant attention and engagement. We pre-registered an exclusion of participants who spend a total of at least 60 seconds off-task throughout the duration of the experiment. Participant engagement was tracked using TaskMaster, which provides measures of time spent on- and off-task (Permut et al., 2019). As per the pre-registration, 266 participants were excluded from the analysis for spending a total of at least 60 seconds off-task throughout the duration of the experiment, suggesting inattentiveness, resulting in 734 participants included in our

preregistered analyses. We also report results with both more and less stringent exclusion criteria.

Design and Procedure

The factor of primary interest was prior exposure (novel vs. exposed), but the full experimental design was a 2 (prior exposure: novel vs. exposed) x 2 (counterbalanced claim set: set A as novel and set B as exposed vs. set B as novel and set A as exposed) x 3 (dependent measure: objectivity vs. fact-checkability vs. existence of truth) x 2 (scale order: e.g., [1] Objective to [6] Subjective vs. [1] Subjective to [6] Objective), where the first factor (prior exposure) is within-participant and of substantive interest and the latter 3 factors (counterbalanced claim set, dependent measure, and scale order) were between-participant and nuisance factors.

The stimulus set consisted of 24 claims, including 8 factual statements (e.g., “President Barack Obama was born in the United States”), 8 opinions (e.g., “Democracy is the greatest form of government”), and 8 borderline claims (e.g., “The Supreme Court must base its rulings on its understanding of what the U.S. Constitution means in current times”), taken from and classified by the Pew Research Center (Mitchell et al., 2018; Doherty, 2018). All claims are given in Table 1. One factual claim was updated to reflect an updated partisan breakdown of members of congress circa April 2022.¹ These claims encompassed a broad range of topics concerning U.S. current events and political beliefs. The experimental design was drawn from the illusory truth literature (e.g., Fazio et al., 2015). The experiment consisted of three stages.

¹ The original claim was “Republicans currently hold a majority of the seats in both chambers of Congress.” The revised claim was “Democrats currently hold a majority of seats in the House of Representatives.”

Table 1

List of Claims Used as Stimuli Accompanied by Claim-Level Means and Standard Deviations for Ratings of Perceived Objectivity.

Classification	Claim	Perceived Objectivity	
		Mean	SD
<i>Facts</i>	Health care costs per person in the U.S. are the highest in the developed world.	5.38	1.15
	President Barack Obama was born in the United States.	5.48	1.22
	Immigrants who are in the U.S. illegally have some rights under the Constitution.	4.72	1.60
	ISIS lost a significant portion of its territory in Iraq and Syria in 2017.	5.17	1.18
	Spending on Social Security, Medicare, and Medicaid make up the largest portion of the U.S. federal budget.	5.28	1.26
	Democrats currently hold a majority of seats in the House of Representatives. ^A	5.58	1.09
	Most of the heroin that currently makes it into the U.S. comes across the southern border.	4.75	1.39
	In the aftermath of the war in Iraq, no active weapons of mass destruction were found.	5.19	1.33
<i>Opinions</i>	Democracy is the greatest form of government.	2.42	1.67
	Increasing the federal minimum wage to \$15 an hour is essential for the health of the U.S. economy.	3.30	1.64
	Abortion should be legal in most cases.	2.40	1.71
	Immigrants who are in the U.S. illegally are a very big problem for the country today.	2.95	1.70
	Government is almost always wasteful and inefficient.	2.87	1.66
	In general, regardless of who is in power, politicians can't be trusted.	2.45	1.62
	The government must make a greater effort to reduce climate change.	3.14	1.81

	The courts have gone too far in restricting public expression of Christian beliefs.	2.55	1.66
<i>Borderline Claims</i>	Applying additional scrutiny to Muslim Americans would not reduce terrorism in the U.S.	3.08	1.62
	Voter fraud across the U.S. has undermined the results of our elections.	3.68	1.88
	Recent tax cuts have benefited the wealthiest of Americans more than others.	4.66	1.49
	Police around the country treat racial and ethnic minorities as fairly as they treat whites.	3.31	1.75
	In the United States, racial discrimination may make it harder for black people to get ahead.	3.63	1.66
	The Supreme Court must base its rulings on its understanding of what the U.S. Constitution means in current times.	3.55	1.72
	Sexual misconduct by Catholic priests and bishops reflects an ongoing problem.	4.33	1.49
	The affordability of a college education is a big problem for the United States.	4.02	1.63

^A Claim was adjusted from Pew to reflect an updated partisan breakdown of members of congress circa April 2022.

NOTE: Classification refers to classification of claims in surveys by the Pew Research Center (Mitchell et al., 2018; Doherty, 2018).

The first stage was an initial exposure stage. Participants were presented with half of the claims in the stimulus set (4 factual statements, 4 opinions, and 4 borderline claims) in three sequential engagement tasks. Participants were asked to: (i) rate the claims for relevance to current events on a 4-item scale from “Not at all relevant” to “Very relevant”; (ii) rate the claims for interestingness on a 4-item scale from “Not at all interesting” to “Very interesting”; (iii) categorize the claims by topic as either “World Events”, “U.S. Events”, “Politics”, “Economics”, “Science”, or “Daily Life”. As a result, by the end of the first

stage participants had seen and engaged with each of the 12 presented claims three times. The particular subset of 12 claims presented in the initial exposure stage was counterbalanced across participants (set A as novel and set B as exposed *or* set B as novel and set A as exposed).

The second stage was a filler stage. The primary purpose of this stage was to separate the first (exposure) stage from the third (classification) stage. During this filler stage, participants completed measures of basic demographics, cognitive reflection using two versions of the cognitive reflection test (Frederick, 2005; Thomson & Oppenheimer, 2016), digital savviness, trust of national news organizations, and political affiliation, adapted from survey measures used by the Pew Research Center (Mitchell et al., 2018).

The third stage was a classification stage, providing our key measures of perceived objectivity. Participants were presented with the full set of 24 claims one at a time. 12 of the claims were repeated from the initial exposure stage and 12 of the claims were novel. Perceived objectivity was elicited using three measures. Participants were randomized to one of the three measures of perceived objectivity (i.e., each participant encountered a single operationalized measure, not all three measures). Using three different measures enabled us to test for robustness across alternative operationalizations of perceived objectivity; we did not anticipate any differences across measures. Each measure consisted of an instruction page with examples and an ascending 6-point scale. These are provided in Figure 1.

Figure 1

Participant Instructions for Measures of Perceived Objectivity

Measure	Instructions
Objectivity	<p>You will now be shown a series of 24 statements. We'd like you to provide a rating for each statement based on the following instructions:</p> <p>Some statements are completely objective, they are verifiable and can be proved or disproved by evidence. For instance "it is raining" is a completely objective statement that is either true or false.</p> <p>Other statements are completely subjective, they are based on the values and beliefs of</p>

the source. For instance “the weather is nice” is a completely subjective statement that one may agree or disagree with but that cannot be definitely proved or disproved.

Many statements contain both objective and subjective information. For instance “it is raining, which is really nice” contains both information that is objective (“it is raining”) and information that is subjective (“which is really nice”).

For each of the statements you will be shown, please rate it on a 6-point scale from [1 (subjective) to 6 (objective) / 1 (objective) to 6 (subjective)]. If you feel strongly that the statement primarily conveys [subjective/objective] information, you should rate it as a 1. If you feel strongly that the statement primarily conveys [objective/subjective] information, you should rate it as a 6.

Please use the entire 6-point scale as feels appropriate.

Fact-checkability	<p>You will now be shown a series of 24 statements. We’d like you to provide a rating for each statement based on the following instructions:</p> <p><i>Snopes</i> is a fact-checking website that uses objective evidence to identify whether various claims are true or false.</p> <p>Fact-checking uses objective evidence to determine the accuracy of a statement. In deciding which claims they can check, <i>Snopes</i> considers whether or not a statement is rooted in a fact that is verifiably true or false. For instance, opinions can’t be fact-checked.</p> <p>On the following pages, imagine that you work for <i>Snopes</i> and that you have been asked to help identify the next batch of claims for fact-checking.</p> <p>For each of the statements you will be shown, please rate it on a 6-point scale based on how possible it is to fact-check from [1 (Definitely cannot be fact-checked) to 6 (Definitely can be fact-checked) / 1 (Definitely can be fact-checked) to 6 (Definitely cannot be fact-checked)]</p> <p>Please use the entire 6-point scale as feels appropriate.</p>
Existence of truth	<p>You will now be shown a series of 24 statements. We’d like you to provide a rating for each statement based on the following instructions:</p> <p>When considering claims, we may ask, <i>Is this claim true?</i> Sometimes there is a correct answer to this question. For instance:</p> <p>“The 2028 Olympic Games will be held in Los Angeles”</p> <p><i>Is this claim true?</i> <u>There is a correct answer to this question</u>, the answer is yes.</p>

“Dusseldorf is the capital of Germany”

Is this claim true? There is a correct answer to this question, the answer is no.

Other times there is no correct answer to this question. For instance:

“Vanilla is the best ice cream flavor”

Is this claim true? There is no correct answer to this question because the claim reflects an opinion.

“Rock music is better than Jazz music”

Is this claim true? There is no correct answer to this question because this claim reflects an opinion.

On the following pages you will be asked, does there exist a correct answer to the question *Is this claim true?* Note, you are not being asked if the claim is true. You are being asked: **Does there exist a correct answer to the question, "Is this claim true?"** For each of the statements you will be shown, please rate it on a 6-point scale from [1 (No correct answer) to 6 (Definitely a correct answer) / 1 (Definitely a correct answer) to 6 (No correct answer)].

Please use the entire 6-point scale as feels appropriate.

The first measure of perceived objectivity, encountered by one third of participants, provided instructions about the notion of objectivity as a construct, along with examples of objective and subjective claims. Participants were then asked to consider each claim and rate it on a numbered 6-point scale from [1] Subjective to [6] Objective. For all measures, scale order was randomized across participants (e.g., for this first measure of perceived objectivity, half of participant saw a scale from [1] Objective to [6] Subjective and half of participants saw a scale from [1] Subjective to [6] Objective). This measure of perceived objectivity was taken verbatim from linguistics research on objectivity (Kaiser & Wang, 2020).

A second measure of perceived objectivity, encountered by another third of participants, provided information about the existence of fact-checking organizations and informed participants that while some claims could be fact-checked, other claims could not be fact-checked, along with examples. Participants were then asked to consider whether it would be possible to have each statement fact-checked by Snopes, a fact-checking organization. Participants were prompted, “Can this statement be fact-checked?” on a

scale from [1] Definitely cannot be fact-checked to [6] Definitely can be fact-checked (scale order was randomized across participants).

A third measure of perceived objectivity instructed another third of participants to consider that when we encounter claims we sometimes wonder about their truthfulness, and so we might ask – is this claim true? While in some instances there is a correct answer to the question about a claim’s truthfulness (i.e., when the claim is true or when the claim is false), in other instances there is no correct answer to the question about a claim’s truthfulness (i.e., when the claim is an opinion). Following the instructions with examples, participants were presented with claims and asked, “Does there exist a correct answer to the question, ‘Is this statement true?’ for this statement?” on a scale from [1] No correct answer to [6] Definitely a correct answer (with scale order randomized across participants). This measure was adapted from Goodwin and Darley (2012) and based on similar measures used in prior work (e.g., Goodwin & Darley, 2008).

In addition to the primary measures of perceived objectivity, at the end of each experiment participants were asked about their agreement with each claim. Presented with the full set of 24 claims (the “exposed” half of which participants were seeing for a fifth time and the “novel” half of which they were seeing for a second time), participants were asked, “To what extent do you agree with the statement below?” on a scale from (1) Strongly disagree to (6) Strongly agree. This scale’s order was never reversed.

Results

The key test of interest for this experiment was the effect of prior exposure on perceived objectivity. Results of analyses are presented as follows. First, we present main effect results collapsed across the three measures of perceived objectivity (as pre-registered) as well as sources of variance for this effect. Second, we explore how the main effect changes based on participant attentiveness, introducing increasingly restrictive attentiveness filtering. Third, we discuss the frequency of the observed main effect across participants. Finally, we analyze the effect of prior exposure on agreement with claims,

and discuss a correlation between the effect of exposure on agreement and the effect of exposure on perceived objectivity .

Main Effect of Prior Exposure on the Perceived Objectivity of Claims

In the experiment, participants were presented with claims and asked to provide ratings of perceived objectivity. Regardless of scale ordering, ratings of perceived objectivity were recorded such that greater values corresponded to greater perceived objectivity. For every individual, we calculate the average rating of perceived objectivity for exposed claims and the average rating of perceived objectivity for novel claims, and take their difference. This difference reflects the effect of exposure on perceived objectivity. We regressed that difference on a contrast coded variable representing the counterbalanced assignment of claim set (1 = set A as novel and set B as exposed, -1 = set B as novel and set A as exposed), two contrast coded variables representing the counterbalanced assignment of a dependent measure of perceived objectivity (-2 = objectivity, 1 = fact-checkability, 1 = existence of truth, and 0 = objectivity, 1 = fact-checkability, -1 = existence of truth, respectively), a contrast coded variable representing the counterbalanced assignment of scale order (1 = objectivity high, -1 = subjectivity high), and all two- and three-way interactions. This pre-registered analysis is equivalent to a mixed ANOVA in which we focus on the main effect of, and interactions with, prior exposure on perceived objectivity. In particular, the intercept represents the main effect of prior exposure, averaged across the 12 between-participant groups, and the various contrast coded variables represent the interactions of those factors with prior exposure. We are most interested in the main effect of prior exposure (the intercept), and consider the other factors as nuisance variables.

Participants perceived exposed claims to be significantly more objective than novel claims ($b = 0.062$, $t(722) = 3.34$, $p < 0.001$, Cohen's $d = 0.123$). Participants' average rating of perceived objectivity was 3.94 for exposed claims and 3.88 for novel claims. This effect is roughly equivalent to that of prior exposure increasing perceived objectivity by one scale point for one of the exposed claims. Table 2 presents a breakdown of the main effect and includes differences in perceived objectivity for exposed vs.

novel claims across different measures of perceived objectivity and different claim types.²

Table 2

Previously Exposed Claims Were Classified as More Objective Than Novel Claims

Measure	Claims	Perceived Objectivity		<i>b</i>	<i>SE</i>	95% CI		<i>df</i>	<i>t</i>	<i>p</i>	<i>Cohen's d</i>
		<i>M</i> _{novel}	<i>M</i> _{exposed}			<i>LL</i>	<i>UL</i>				
Across measures of perceived objectivity	All Claims	3.88	3.94	0.062	0.019	0.026	0.099	722	3.34	<0.001	0.123
	Facts	5.16	5.22	0.064	0.026	0.013	0.115	722	2.46	0.014	
	Opinions	2.71	2.81	0.101	0.035	0.033	0.169	722	2.90	0.004	
	Borderline	3.77	3.79	0.021	0.037	-0.051	0.094	722	0.58	0.562	
Existence of truth	All Claims	3.98	4.05	0.067	0.034	<0.001	0.134	234	1.98	0.048	0.129
	Facts	5.19	5.26	0.065	0.047	-0.027	0.156	234	1.39	0.166	
	Opinions	2.85	2.94	0.097	0.065	-0.031	0.225	234	1.49	0.137	
	Borderline	3.91	3.95	0.041	0.065	-0.087	0.168	234	0.63	0.529	
Objectivity	All Claims	3.55	3.60	0.053	0.034	-0.014	0.119	237	1.57	0.118	0.101
	Facts	4.94	5.04	0.089	0.052	-0.012	0.192	237	1.74	0.083	
	Opinions	2.37	2.44	0.079	0.056	-0.032	0.190	237	1.40	0.162	
	Borderline	3.33	3.32	-0.011	0.069	-0.147	0.125	237	-0.15	0.878	
Fact-checking	All Claims	4.10	4.17	0.066	0.029	0.009	0.124	251	2.27	0.024	0.142
	Facts	5.33	5.37	0.038	0.037	-0.035	0.111	251	1.03	0.304	
	Opinions	2.91	3.03	0.127	0.060	0.010	0.244	251	2.13	0.034	
	Borderline	4.07	4.10	0.034	0.059	-0.082	0.150	251	0.58	0.561	

The omnibus F test was not statistically significant ($F(11, 722) = 1.10, p = 0.356$), suggesting that, overall, the main effect did not vary across the 12 groups. More specifically, there was no interaction effect of prior exposure with the particular dependent measure of perceived objectivity that was used ($F(2, 722) = 0.06, p = 0.938$) nor with the direction of measurement scale order ($F(1, 722) = 0.09, p =$

² One might instead classify each response as “objective” or “subjective” by splitting the response scale at its midpoint. Using such an approach and aggregating across stimuli, we find that 60.3% of exposed claims are classified as objective and 58.9% of novel claims are classified as objective, a statistically significant difference ($b = 1.39, t(722) = 2.44, p = 0.015$).

0.768). The counterbalanced assignment of claim subset as exposed or as novel was statistically significant ($F(1, 722) = 4.88, p = 0.028$), indicating some potential evidence for heterogeneity of the effect across claims. None of the two- or three-way interactions among the three factors were significant (all $ps > 0.120$) The results above constitute our pre-registered analyses. Table 3 presents the complete ANOVA results. The main effect of prior exposure on perceived objectivity was not correlated with any of the measured individual differences such as performance on the cognitive reflection test, education, political liberalism, gender, etc.

Table 3

Complete ANOVA of Main Effect of Prior Exposure on Perceived Objectivity

	Df	Sum Sq	Mean Sq	<i>F</i>	<i>p</i>
Intercept (Main Effect of Prior Exposure)	1	2.83	2.83	11.14	<0.001
Claim Set	1	1.24	1.24	4.88	0.028
Scale Order	1	0.02	0.02	0.09	0.768
Perceived Objectivity Measure	2	0.03	0.02	0.06	0.938
Claim Set × Scale Order	1	0.13	0.13	0.53	0.468
Claim Set × Perceived Objectivity Measure	2	1.08	0.54	2.12	0.120
Scale Order × Perceived Objectivity Measure	2	0.03	0.01	0.05	0.947
Claim Set × Scale Order × Perceived Objectivity Measure	2	0.51	0.26	1.00	0.368
Residuals	722	183.49	0.25	-	-

Given that we are interested in the effect of prior exposure on the perceived objectivity of claims in general, rather than just on the specific claims used in this experiment, there is value in considering an analysis that generalizes across stimuli as well as individuals (e.g., Judd, Westfall, & Kenny, 2012). An additional exploratory analysis used a mixed-model regression with perceived objectivity ratings of each

claim by each participant as the dependent variable and treated individual claims as random factors. Perceived objectivity was regressed on prior exposure (1 = exposed, -1 = novel), the counterbalanced assignment of a dependent measure of perceived objectivity (-2 = objectivity, 1 = fact-checkability, 1 = existence of truth, and 0 = objectivity, 1 = fact-checkability, -1 = existence of truth, respectively), the counterbalanced assignment of claim set (1 = set A as novel and set B as exposed, -1 = set B as novel and set A as exposed), the counterbalanced assignment of scale order (1 = objectivity high, -1 = subjectivity high), and all interactions among different factors. The main effect of exposure on perceived objectivity in the mixed-model analysis was consistent with the main effect observed using pre-registered analyses ($b = 0.062$, $t(17,616) = 2.63$, $p = 0.011$) with relatively little variance in the main effect size across claims ($SD = 0.023$) or people ($SD = 0.005$). This suggests that the magnitude of the effect is relatively consistent across the 24 claims, rather than being simply driven by a subset of unusual claims.

Exploring the Main Effect Across More/Less Attentive Participants

The integrated use of TaskMaster (Permut et al., 2019) throughout the experiment allowed for us to track participant attentiveness and engagement. For every page of the web survey, TaskMaster tracks the amount of time (in seconds) that participants spend on-task (on the webpage) and off-task (off the webpage). The primary analysis of the main effect presented earlier followed the pre-registered exclusion of all participants who spent a total of at least 60 seconds off-task throughout the duration of the experiment. But it is also possible to consider what analyses look like when using other attention cutoffs, either more or less restrictive ones.

First, we consider the case of no attention filtering. When repeating primary analyses with the inclusion of all participants who completed the experiment, we find that perceived objectivity is greater for exposed claims than for novel claims ($b = 0.045$, $t(988) = 2.89$, $p = 0.004$, Cohen's $d = 0.091$).

Next, we consider more restrictive attention filters. Participants are divided into three groups: (i) participants who spent a total of at least 1 minute off-task throughout the duration of the experiment ($N = 266$); (ii) participants who spent a total of more than 0 but less than 1 minute off-task throughout the

duration of the experiment ($N = 392$); (iii) participants who spent no time at all off-task throughout the duration of the experiment ($N = 342$). In the primary analyses, we pre-registered the exclusion of participants in group (i) and analyzed across participants in groups (ii) and (iii). Here, we consider the effect for each group separately. We do not find a main effect for participants in segment (i) who spent at least 1 minute off-task ($b = 0.002$, $t(254) = 0.07$, $p = 0.944$). For participants in segment (ii) who spent more than 0 but less than 1 minute off-task, we find a non-significant effect in the expected direction ($b = 0.040$, $t(380) = 1.57$, $p = 0.118$). Using the most restrictive attention filter, for participants in segment (iii) who did not spend any time at all off-task, we find a significant main effect that is directionally greater than the main effect from the primary analyses reported earlier ($b = 0.094$, $t(330) = 3.36$, $p < 0.001$, Cohen's $d = 0.182$). Taken together, analyses of attentiveness data suggest that more rigorously filtering out inattentive participants who spend time off-task results in a greater effect of prior exposure on perceived objectivity.³

Frequency of the Creeping Objectivity Effect Among Participants

In addition to the magnitude of the effect size, we are interested in how frequently the observed effect occurs. In other words, what percentage of participants have higher ratings of perceived objectivity for exposed claims than for novel claims? We find that 51% of participants report greater perceived objectivity for exposed claims than for novel claims, 40% of participants report greater perceived objectivity for novel claims than for exposed claims, and 9% of participants report no difference in perceived objectivity between exposed claims and novel claims. If employing more restrictive filtration criteria and only including participants who were on-task for the duration of the experiment, more participants exhibit the effect, with 55%, 37%, and 8% exhibiting the expected effect, the reversed effect, and no effect, respectively.

Effect of Prior Exposure on Agreement with Claims

³ If instead one classifies claims using the categorical measure described in footnote 2, the effect of prior exposure on the percentage of claims classified as objective for each segment would be (i) ($b = -0.60$, $t(254) = -0.68$, $p = 0.495$); (ii) ($b = 0.95$, $t(380) = 1.22$, $p = 0.223$); (iii) ($b = 2.01$, $t(330) = 2.40$, $p = 0.017$).

To assess the impact of exposure on agreement with claims, for every individual we calculate the average rating of agreement with exposed claims and the average rating of agreement with novel claims, and take their difference. This difference reflects the effect of perceived objectivity on agreement. We regressed that difference on the same set of contrast coded variables used for our primary preregistered analysis. This pre-registered analysis is equivalent to a mixed ANOVA in which we focus on the main effect of, and interactions with, prior exposure on agreement. In particular, the intercept represents the main effect of prior exposure, averaged across the 12 between-participant groups, and the various contrast coded variables represent the interactions of those factors with prior exposure. We are most interested in the main effect of prior exposure (the intercept), and consider the other factors as nuisance variables.

On average, agreement with exposed claims was 4.30 and with novel claims was 4.27. We find that there was no significant difference in agreement with exposed vs. novel claims ($b = 0.025$, $t(722) = 1.53$, $p = 0.126$), nor any significant variance in this null effect across counterbalanced groups ($F(11, 722) = 0.88$, $p = 0.562$). However, agreement was significantly greater with exposed opinions than with novel opinions ($b = 0.071$, $t(722) = 2.25$, $p = 0.025$).

Finally, we also consider the correlation between the effect of exposure on agreement and the effect of exposure on perceived objectivity. For every participant, there is a key measure reflecting the effect of exposure on perceived objectivity (the difference in perceived objectivity for exposed vs. novel claims) and a secondary measure reflecting the effect of exposure on agreement (the difference in agreement with exposed vs. novel claims). We regressed the key measure of perceived objectivity for exposed vs. novel claims on an intercept, a complete set of 11 contrast coded variables, and on the secondary measure of agreement with exposed vs. novel claims. We find that the difference in perceived objectivity of exposed vs. novel claims is significantly positively correlated with the difference in agreement with exposed vs. novel claims ($b = 0.269$, $t(721) = 6.49$, $p < 0.001$). This result reflects a significant positive correlation between the main effect of prior exposure on perceived objectivity and the effect of prior exposure on agreement.

General Discussion

In a three-stage experiment, we explored the effect of prior exposure on perceived objectivity. Across a variety of claims and three different measures of perceived objectivity, we find that previously exposed claims are rated as more objective than are novel claims. On average, across participants, claims, and measures, prior exposure increases perceived objectivity for a claim by approximately 0.06 scale points and by approximately 0.09 scale points among participants who show the least evidence of inattention. Moreover, the average rating of perceived objectivity was greater for exposed than for novel claims for 51% of participants, with only 40% showing an effect in the opposite direction.

When considering useful benchmarks to understand effect size, it is possible to look to the well-documented illusory truth effect, a related, albeit distinct paradigm documenting the effect of exposure on the perceived veracity of factual statements. For instance, in Fazio et al., (2019) on average 48% of novel facts were classified as true and 52% of exposed facts were classified as true. The magnitude of the illusory truth effect varied by claim plausibility, such that highly plausible and highly implausible claims had a reduced effect (Fazio et al., 2019). In the current research, the effect of prior exposure on perceived objectivity was studied using claims about political and social issues largely taken verbatim from Pew Research Center surveys (Mitchell et al., 2018; Doherty, 2018). In a meta-analysis of the illusory truth effect, Henderson et al. (2021) find that a majority of the prior research on the illusory truth effect used stimuli comprised of trivia claims, likely in order to ensure the true novelty of ‘novel’ claims, rather than using claims of high day-to-day visibility. Recent work on the illusory truth effect for news headlines used claims more central to ongoing social and political discourse, while also aiming to ensure claim novelty by asking participants whether or not they had seen the claims previously (e.g., Pennycook et al., 2018). De Keersmaecker et al. (2020), replicating the findings of Pennycook et al. (2018), found that repeated exposure of a new headline increases perceived accuracy on average by 0.09 scale points on a 4-point scale.

It is also likely that the magnitude of this creeping objectivity effect varies across contexts outside

of the current experimental design. We find mixed evidence regarding variation across specific claims: the effect was larger for one set of 12 claims than the other, but a model treating claims as random effects indicated relatively little variability across claims. The stimulus set used in the current design is comprised of claims taken directly from Pew Research Center surveys (i.e., selected for downstream relevance and interest to participants but not cherry-picked for selectively larger effect sizes compared with other claims). These claims are central to a variety of public opinions for American respondents. People may have strong beliefs about the validity and objectivity of these claims, making it more difficult to detect shifts in beliefs than for other claims (e.g., for novel ideas or emerging topics of discussion). Relatedly, it is quite possible or even likely that many participants are not seeing these claims for the first time, even when presented in our experimental design as ‘novel’. Hence, it may be that the detected effect of prior exposure on perceived objectivity is in fact the difference between an n^{th} and an $n+3^{\text{rd}}$ exposure rather than a 1st vs. a 4th exposure. Although we don’t have a metric of how frequently participants had been exposed to particular claims outside of the experimental setting, if an attenuated effect of exposure on perceived objectivity was observed under such circumstances it would be in line with findings of attenuated strength for the illusory truth effect in instances of subsequent versus initial repetition (Hassan and Barber, 2021). These results, along with the finding that more attentive participants exhibit a bigger effect, imply that there likely exists variation across people, exposure contexts, and individual claims, for which an effect of exposure on perceived objectivity may be stronger than the modest effect observed in this particular setting.

Another question of interest is whether or not participants are accurately identifying facts as objective claims and opinions as subjective claims. Although measures of perceived objectivity used 6-point scales, we can bisect these scales to categorize responses as either “subjective” or as “objective”. Accurate claim classification can then be defined as the percentage of facts a participant classifies as “objective” and opinions as “subjective”, based on the a-priori classification of claims by the Pew Research Center. On average, we find that 77% of exposed claims were accurately classified and 78% of

novel claims were accurately classified, with no significant effect of prior exposure on accurate claim classification ($b = -0.86$, $t(722) = -1.38$, $p = 0.169$). However, prior exposure did result in significantly decreased classification accuracy for opinions ($b = -3.13$, $t(722) = -3.14$, $p = 0.002$), likely an outcome of the main effect of creeping objectivity. Accurate claim classification was also positively correlated with individual differences such as better performance on the cognitive reflection test ($b = 3.46$, $t(721) = 9.75$, $p < 0.001$) and political liberalism ($b = 2.49$, $t(718) = 3.45$, $p < 0.001$).

In addition to measures of perceived objectivity, participants were also asked to provide ratings of agreement with claims. As measures of agreement were collected following the third (classification) stage of the experiment, participants saw claims they were rating on agreement for either a 2nd or a 5th time. It is possible that diminishing marginal effects limited our ability to detect an effect of exposure on agreement with factual claims. The measure of agreement was intended to track general endorsement of claims regardless of how objective they were perceived, but it is possible that this language proved more jarring for endorsement of factual claims and prevented our ability to detect an exposure effect on agreement with factual claims. While we did not see an effect of prior exposure on overall agreement, we did find that prior exposure increases agreement with opinions, consistent with the mere exposure effect (Cacioppo & Petty, 1979).

Across claims and measures, the difference in agreement with exposed vs. novel claims was significantly correlated with the difference in perceived objectivity of exposed vs. novel claims. Both across claims, and when examining factual claims and opinions separately, we find that individuals who show a larger difference between exposed and novel claims in agreement also show a larger difference between exposed and novel claims in perceived objectivity. However, this correlation did not entirely account for the effect of exposure on perceived objectivity.

Process evidence

Why does prior exposure affect beliefs about the objectivity of claims? It is possible that prior exposure enhances endorsement (agreement or perceived accuracy) of claims and, in turn, increased claim

endorsement enhances the belief that a claim is more objective (in line with naïve realism). Our results are suggestive of this possible causal chain, though are not yet conclusive. In additional experiments, described in brief below with further details in supplemental materials, we explored processing fluency and perceived social consensus as potential mechanisms. However, these additional experiments do not allow us to make strong inferences regarding potential processes of the creeping objectivity effect.

Experiment C1 used a non-repetition-based visual manipulation to investigate processing fluency as a potential process for the observed effect of prior exposure on perceived objectivity, given prior findings that processing fluency at least partially explains the illusory truth and mere exposure effects (Oppenheimer, 2006; Shah and Oppenheimer, 2007; Alter and Oppenheimer, 2009). Participants were presented with a subset of the claims used in the three-stage experiment described above. Half of the claims were presented in a visually fluent manner (high text-to-background contrast) and half of the claims were presented in a visually disfluent manner (low text-to-background contrast), randomized across participants. During that initial presentation, participants were asked to classify each claim as either fact-checkable or as not fact-checkable. We find that a greater percentage of fluent opinions were classified as fact-checkable (30%) than disfluent opinions were classified as fact-checkable (27%) representing a significant difference in the perceived objectivity for fluent vs. disfluent opinions ($b = 3.31$; $t(330) = 2.55$, $p = 0.011$). However, there is no such significant difference in perceived objectivity between fluent and disfluent factual claims ($b = -0.38$; $t(330) = -0.33$, $p = 0.744$). On average, 88% of factual claims were classified as fact-checkable, so it is possible that ceiling effects and the use of a binary measure of perceived objectivity prevented us from being able to detect an effect of visual fluency on perceived objectivity of factual claims. When collapsing across claim types, we do not find that visually fluent claims are any more likely to be classified as fact-checkable than visually disfluent claims ($b = 0.813$; $t(330) = 0.91$, $p = 0.365$). Although a substantial body of literature has demonstrated the importance of processing fluency for the illusory truth and mere exposure effects, and we find some suggestive evidence based on an effect of fluency for opinions, we are unable to conclude from the results of this experiment

that processing fluency is a key mechanism of the observed effect of repeated exposure on perceived objectivity. Future investigations into the role of fluency on the observed creeping objectivity effect may merit alternative manipulations of processing fluency as well as other related processes such as answer fluency (Thompson et al., 2013).

However, the effects of repetition may not be limited to fluency. An illusory truth effect is observed even when people are simply *told* that a claim has been repeated, without being repeatedly exposed to it themselves, suggesting that people believe repetition is itself an informative cue about the nature of a claim (Mattavelli et al., 2022). Experiment C2 investigated perceived social consensus as a potential process for the observed effect of prior exposure on perceived objectivity. More widely endorsed beliefs are perceived as more objective (Young & Heiphetz, 2007; Goodwin and Darley, 2008, 2012) and prior exposure to a claim increases its perceived social consensus (Weaver et al., 2007). Together, these findings raise the possibility that prior exposure may increase perceived objectivity by increasing the perceived social consensus around a claim. Experiment C2 used the same subset of claims as was used in experiment C1 and manipulated exposure to half of the claims using a three-stage experimental design similar to that described above. In the third stage, rather than collecting measures of perceived objectivity, participants were asked to report the percentage of U.S. citizens they think would agree with each claim. Half of the claims for which participants provided ratings of perceived social consensus had been exposed once in the first stage of the experiment and half of the claims were novel, randomized across participants. Across participants, the average perceived social consensus for exposed claims was 49.68% and the average perceived social consensus for novel claims was 49.54%. Whether or not claims had been previously exposed did not affect their perceived social consensus ($b = 0.141$, $t(386) = 0.31$, $p = 0.756$). Separating by claim type does not yield additional insight. Results from prior research notwithstanding, we fail to find an effect of prior exposure on perceived social consensus and hence are unable to conclude whether or not perceived social consensus may mediate the effect of prior exposure on perceived objectivity.

Future Directions and Limitations

The effect of prior exposure found in the current research highlights that the manner in which a claim is presented influences beliefs about the claim's objectivity. Two potentially insightful avenues for future research stem from possible changes to the current experimental design. The first is to consider different exposure frequencies or intensities. The three-stage design outlined above manipulated prior exposure by presenting a subset of claims in the first stage, inducing a brief delay during the second (filler) stage, and measuring the perceived objectivity of the full set of claims in the third stage. During the first stage, exposed claims were each presented to participants three times in three different engagement tasks. The second direction is to experiment with varying filler stage delay periods. On average, participants spent 169 seconds on the questions in the filler stage, which included measures of basic demographics, an extended version of the cognitive reflection test, as well as measures of digital savviness, trust of national news organizations, political awareness, etc., each adapted from measures used in Pew Research Center surveys. Either of these directions could lend considerable insight into the creeping objectivity effect.

In the current research we found that, across a variety of claims and measures, prior exposure increases the perceived objectivity of a claim. This creeping objectivity effect was present for a majority of participants. While the observed effect was modest in magnitude, such pervasive effects can have substantial implications on large scales and for complex problems. The creeping objectivity effect is of note to and may merit further attention from researchers of the illusory truth and mere exposure effects as well as consumer behavior, conflict mediation, and misinformation more broadly.

Illusory Truth and Mere Exposure Research

One potential methodological contribution of the current findings may be of note for research on the illusory truth and mere exposure effects. Prior research on both the illusory truth and the mere exposure effects has often presented participants with claims that are a priori determined to be either factual statements or opinions. As a result, elicitation of manipulations intended to shift perceptions of

accuracy or agreement may be unable to detect shifts in perceived objectivity. For example, repeated exposure to an opinion may result in a fraction of participants starting to perceive it as a fact, but dependent measures (e.g., agreement) may not reflect such a shift in perceived objectivity. This methodological nuance is important for researchers, as it could affect how responses are coded and results are interpreted. The present findings indicate that not only does prior exposure affect whether individuals believe factual claims are true, or opinions are agreeable, but prior exposure can also affect whether claims are perceived as matters of objective fact or of subjective opinion. Future research using these paradigms may wish to consider including measures of perceived objectivity in addition to a-priori determined measures of agreement or accuracy.

Consumer Behavior

For consumer researchers and marketing managers, the current findings present implications in the use of claims comparing products or describing product attributes. For instance, comparisons between brands are often presented to consumers as direct superiority claims about product quality. Consumers are willing to pay more for the superior product in question when they believe that the comparison is being made along a dimension of objective quality rather than along a dimension of subjective taste (Spiller & Belogolova, 2016). The current research suggests that prior exposure to such claims may increase the likelihood that claims are perceived as reflecting objective dimensions such as quality, rather than subjective dimensions such as taste.

Conflict Management

Moral objectivism and naïve realism are focal issues in work on interpersonal conflict (Ross & Ward, 1995). The belief that one's views are reflective of an objective state where there is no room for disagreement limits people's receptiveness to alternative views (Lieberman et al., 2012; Minson et al., 2020; Yeomans et al., 2020). This results in workplace inefficiencies through hindered collaboration as well as increased political polarization through belief entrenchment and moral tribalism (Lieberman et al., 2012; Skitka & Morgan, 2014). The present findings suggest that an increase in the perceived objectivity

of repeatedly exposed claims may be associated with some of the divisive features of interpersonal conflict and with the downstream consequence of societal fractures. Future research would benefit from directly investigating these potential associations causally.

Misinformation

With respect to the rise of misinformation disseminated via social media, social media companies have come under pressure to not only fact-check factual statements but also to discern factual statements from opinions (Iannucci, 2017; Media Insight Project, 2018). For instance, Twitter added a notice reading “Get the facts about mail-in ballots” to two of President Trump’s tweets, indicating that the tweets included fact-checkable claims that could be (and were) false (Conger & Isaac, 2020). Non-governmental organizations such as the International Research & Exchanges Board (IREX) have developed international programs dedicated to teaching both journalists and laypersons to discern fake news from real news (Murrock et al., 2018). Our findings suggest that initiatives targeted at improving the accuracy of information consumption could also benefit from teaching people to discern news from opinions.

Conclusion

As communicators, consumers, and colleagues, we encounter claims on a daily basis about issues in the world around us, some of which are objective and are either correct or incorrect, and others of which are subjective and reflect varying views and opinions. Such perceived objectivity has downstream consequences on collaboration, conflict, consumer behavior, political polarization, and misinformation. Moreover, we frequently encounter the same claims multiple times. Sometimes repeated exposure to claims is a result of issues becoming focal points of discussion, debate, or disagreement, and claims reflecting those issues are repeated by our friends, coworkers, and media outlets. Other times repeated exposure to claims may be the result of marketing or political campaigns or popularized slogans. Whatever the source of initial exposure, repeated exposure to a claim affects the way in which it is processed and understood. Across a variety of claims and measures, we uncover a persistent effect: prior exposure to a claim increases the claim’s perceived objectivity. The current research bridges prior

findings in the illusory truth and mere exposure literatures, reinforcing the notion that whether a claim is perceived as objective or as subjective is itself a malleable construct.

References

- Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review, 13*(3), 219-235.
- Arkes, H. R., Hackett, C., & Boehm, L. (1989). The generality of the relation between familiarity and judged validity. *Journal of Behavioral Decision Making, 2*(2), 81-94.
- Arkes, H. R., Boehm, L. E., & Xu, G. (1991). Determinants of judged validity. *Journal of Experimental Social Psychology, 27*(6), 576-605.
- Bacon, F. T. (1979). Credibility of repeated statements: Memory for trivia. *Journal of Experimental Psychology: Human Learning and Memory, 5*(3), 241.
- Berryman, J. C. (1984). Interest and liking: Further sequential effects. *Current Psychological Research & Reviews, 3*(4), 39-42.
- Begg, I. M., Anas, A., & Farinacci, S. (1992). Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General, 121*(4), 446.
- Blatz, C. W., & Mercier, B. (2018). False polarization and false moderation: Political opponents overestimate the extremity of each other's ideologies but underestimate each other's certainty. *Social Psychological and Personality Science, 9*(5), 521-529.
- Bornstein, R. F. (1989). Exposure and affect: Overview and meta-analysis of research, 1968–1987. *Psychological Bulletin, 106*(2), 265.
- Cacioppo, J. T., & Petty, R. E. (1979). Effects of message repetition and position on cognitive response, recall, and persuasion. *Journal of Personality and Social Psychology, 37*(1), 97.
- Crandall, C. S. (1985). The liking of foods as a result of exposure: Eating doughnuts in Alaska. *The Journal of Social Psychology, 125*(2), 187-194.
- De Keersmaecker, J., Dunning, D., Pennycook, G., Rand, D. G., Sanchez, C., Unkelbach, C., & Roets, A. (2020). Investigating the robustness of the illusory truth effect across individual differences in

- cognitive ability, need for cognitive closure, and cognitive style. *Personality and Social Psychology Bulletin*, 46(2), 204-215.
- 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(2), 238-257.
- Doherty, C. (2018). Key findings on Americans' views of the US political system and democracy. *Pew Research Center*.
- Effron, D. A., & Raj, M. (2020). Misinformation and morality: encountering fake-news headlines makes them seem less unethical to publish and share. *Psychological Science*, 31(1), 75-87.
- 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(5), 993.
- Fazio, L. K., Rand, D. G., & Pennycook, G. (2019). Repetition increases perceived truth equally for plausible and implausible statements. *Psychonomic Bulletin & Review*, 1-6.
- Fisher, M., Knobe, J., Strickland, B., & Keil, F. C. (2017). The influence of social interaction on intuitions of objectivity and subjectivity. *Cognitive science*, 41(4), 1119-1134.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4), 25-42.
- Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106(3), 1339-1366.
- Goodwin, G. P., & Darley, J. M. (2012). Why are some moral beliefs perceived to be more objective than others?. *Journal of Experimental Social Psychology*, 48(1), 250-256.
- Griffin, D. W., & Ross, L. (1991). Subjective construal, social inference, and human misunderstanding. In *Advances in Experimental Social Psychology* (Vol. 24, pp. 319-359). Academic Press.
- Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior*, 16(1), 107-112.

- Hassan, A., & Barber, S. J. (2021). The effects of repetition frequency on the illusory truth effect. *Cognitive research: principles and implications*, 6(1), 1-12.
- Hawkins, S. A., & Hoch, S. J. (1992). Low-involvement learning: Memory without evaluation. *Journal of Consumer Research*, 19(2), 212-225.
- Heiphetz, L., & Young, L. L. (2017). Can only one person be right? The development of objectivism and social preferences regarding widely shared and controversial moral beliefs. *Cognition*, 167, 78-90.
- Henderson, E. L., Westwood, S. J., & Simons, D. J. (2021). A reproducible systematic map of research on the illusory truth effect. *Psychonomic bulletin & review*, 1-24.
- Hill, W. F. (1978). Effects of mere exposure on preferences in nonhuman mammals. *Psychological Bulletin*, 85(6), 1177.
- Iannucci, R. (2017, August 16). News or opinion? Online, it's hard to tell. Retrieved July 07, 2020, <https://www.poynter.org/ethics-trust/2017/news-or-opinion-online-its-hard-to-tell/>
- Jacoby, L. L., Kelley, C., Brown, J., & Jasechko, J. (1989). Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of Personality and Social Psychology*, 56(3), 326.
- Johar, G. V., & Roggeveen, A. L. (2007). Changing false beliefs from repeated advertising: The role of claim-refutation alignment. *Journal of Consumer Psychology*, 17(2), 118-127.
- Johnson, S. G., Rodrigues, M., & Tuckett, D. (2021). Moral tribalism and its discontents: How intuitive theories of ethics shape consumers' deference to experts. *Journal of Behavioral Decision Making*, 34(1), 47-65.
- Kaiser, E., & Wang, C. (2020). Distinguishing fact from opinion: Effects of linguistic packaging. In S. Denison, M. Mack, Y. Xu, & B. Armstrong (Eds.), *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*, (pp. 116-122), Cognitive Science Society.

- Kaiser, E., & Wang, C. (2021). Packaging information as fact versus opinion: Consequences of the (information-) structural position of subjective adjectives. *Discourse Processes*, 58(7), 617-641.
- Lieberman, V., Minson, J. A., Bryan, C. J., & Ross, L. (2012). Naïve realism and capturing the “wisdom of dyads”. *Journal of Experimental Social Psychology*, 48(2), 507-512.
- Litman, L., Rosenzweig, C., & Moss, A. (2020). New Solutions Dramatically Improve Research Data Quality on MTurk.
- Mattavelli, Simone & Corneille, Olivier & Unkelbach, Christian. (2022). Truth by Repetition ... without repetition: Testing the effect of instructed repetition on truth judgments. *Journal of Experimental Psychology Learning Memory and Cognition*.
- McDougal v. Fox News Network LLC, No. 19-cv-11161-MKV, Docket no. 37 (SDNY June 25, 2020).
- Media Insight Project. (2018). Americans and the news media: What they do—and don’t—understand about each other. *American Press Institute*.
- Minson, J. A., Chen, F. S., & Tinsley, C. H. (2020). Why won’t you listen to me? Measuring receptiveness to opposing views. *Management Science*, 66(7), 3069-3094.
- Mitchell, A., Gottfried, J., & Weisel, R. (2018). Distinguishing Between Factual and Opinion Statements in the News. *Pew Research Center*.
- Murrock, E., Amulya, J., Druckman, M., & Liubyva, T. (2018). Winning the War on State-Sponsored Propaganda: Results from an Impact Study of a Ukrainian News Media and Information Literacy Program. *Journal of Media Literacy Education*, 10(2), 53-85.
- Oppenheimer, D. M. (2006). Consequences of erudite vernacular utilized irrespective of necessity: Problems with using long words needlessly. *Applied Cognitive Psychology*, 20(2), 139-156.
- Penney, V. (2020, July 14). How Facebook Handles Climate Disinformation. Retrieved July 17, 2020, <https://www.nytimes.com/2020/07/14/climate/climate-facebook-fact-checking.html>
- Pennycook, G., Cannon, T. D., & Rand, D. G. (2018). Prior exposure increases perceived accuracy of fake news. *Journal of Experimental Psychology: General*.

- Permut, S., Fisher, M., & Oppenheimer, D. M. (2019). Taskmaster: A tool for determining when subjects are on task. *Advances in Methods and Practices in Psychological Science*, 2(2), 188-196.
- Petrocelli, J. V., Tormala, Z. L., & Rucker, D. D. (2007). Unpacking attitude certainty: Attitude clarity and attitude correctness. *Journal of Personality and Social Psychology*, 92(1), 30.
- Pliner, P. (1982). The effects of mere exposure on liking for edible substances. *Appetite*, 3(3), 283-290.
- Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science*, 9(1), 45-48.
- Ross, L., & Ward, A. (1995). Psychological barriers to dispute resolution. In *Advances in experimental social psychology* (Vol. 27, pp. 255-304). Academic Press.
- Russell, D. W. (1996). UCLA Loneliness Scale (Version 3): Reliability, validity, and factor structure. *Journal of Personality Assessment*, 66(1), 20-40.
- Sarkissian, H., Park, J., Tien, D., Wright, J. C., & Knobe, J. (2011). Folk moral relativism. *Mind & Language*, 26(4), 482-505.
- Shah, A. K., & Oppenheimer, D. M. (2007). Easy does it: The role of fluency in cue weighting. *Judgment and Decision Making*, 2(6), 371-379.
- Shane, S. (2017). From headline to photograph, a fake news masterpiece. *The New York Times*, 18.
- Shidlovski, D., Schul, Y., & Mayo, R. (2014). If I imagine it, then it happened: The Implicit Truth Value of imaginary representations. *Cognition*, 133(3), 517-529.
- Skitka, L. J., & Morgan, G. S. (2014). The social and political implications of moral conviction. *Political psychology*, 35, 95-110.
- Spiller, S. A., & Belogolova, L. (2016). On consumer beliefs about quality and taste. *Journal of Consumer Research*, 43(6), 970-991.
- Thomson, K. S., & Oppenheimer, D. M. (2016). Investigating an alternate form of the cognitive reflection test. *Judgment and Decision Making*, 11(1), 99.

- Thompson, V. A., Turner, J. A. P., Pennycook, G., Ball, L. J., Brack, H., Ophir, Y., & Ackerman, R. (2013). The role of answer fluency and perceptual fluency as metacognitive cues for initiating analytic thinking. *Cognition*, *128*(2), 237-251.
- Toner, K., Leary, M. R., Asher, M. W., & Jongman-Sereno, K. P. (2013). Feeling superior is a bipartisan issue: Extremity (not direction) of political views predicts perceived belief superiority. *Psychological Science*, *24*(12), 2454-2462.
- 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*(1), 219.
- Unkelbach, C., & Rom, S. C. (2017). A referential theory of the repetition-induced truth effect. *Cognition*, *160*, 110-126.
- US Dominion, Inc. v. Powell*, Civil Action 1:21-cv-00040 (CJN) (D.D.C. Aug. 11, 2021)
- Wang, W. C., Brashier, N. M., Wing, E. A., Marsh, E. J., & Cabeza, R. (2016). On known unknowns: Fluency and the neural mechanisms of illusory truth. *Journal of Cognitive Neuroscience*, *28*(5), 739-746.
- Whittlesea, B. W. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*(6), 1235.
- Yeomans, M., Minson, J., Collins, H., Chen, F., & Gino, F. (2020). Conversational receptiveness: Improving engagement with opposing views. *Organizational Behavior and Human Decision Processes*, *160*, 131-148.
- Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, *9*(2p2), 1-27.