

‘You can’t bullshit a bullshitter’ (or can you?):
Bullshitting frequency predicts receptivity to various
types of misleading information

Shane Littrell, Evan F. Risko, & Jonathan A. Fugelsang

Department of Psychology, University of Waterloo, Waterloo, ON, Canada

This is the peer reviewed version of the following article: **Littrell, S., Risko, E. F., & Fugelsang, J. A. (2021). 'You can't bullshit a bullshitter' (or can you?): Bullshitting frequency predicts receptivity to various types of misleading information. *British Journal of Social Psychology*, 60(4), 1484-1505.**, which has been published in final form at <https://doi.org/10.1111/bjso.12447>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions. This article may not be enhanced, enriched or otherwise transformed into a derivative work, without express permission from Wiley or by statutory rights under applicable legislation. Copyright notices must not be removed, obscured or modified. The article must be linked to Wiley's version of record on Wiley Online Library and any embedding, framing or otherwise making available the article or pages thereof by third parties from platforms, services and websites other than Wiley Online Library must be prohibited.

Correspondence concerning this article should be addressed to Shane Littrell, Department of Psychology, University of Waterloo, Waterloo, ON, Canada N2L 3G1

CONTACT: Shane Littrell, sslittrell@uwaterloo.ca

Abstract

Research into both receptivity to falling for bullshit and the propensity to produce it have recently emerged as active, independent areas of inquiry into the spread of misleading information. However, it remains unclear whether those who frequently produce bullshit are inoculated from its influence. For example, both bullshit receptivity and bullshitting frequency are negatively related to cognitive ability and aspects of analytic thinking style, suggesting that those who frequently engage in bullshitting may be *more likely* to fall for bullshit. However, separate research suggests that individuals who frequently engage in deception are better at detecting it, thus leading to the possibility that frequent bullshitters may be *less likely* to fall for bullshit. Here we present three studies ($N = 826$) attempting to distinguish between these competing hypotheses, finding that frequency of *persuasive bullshitting* (i.e., bullshitting intended to impress or persuade others) positively predicts susceptibility to various types of misleading information and that this association is robust to individual differences in cognitive ability and analytic cognitive style.

Key words: bullshitting, bullshit receptivity, fake news, misinformation, deception detection

Introduction

Assessing the cognitive mechanisms underlying the transmission and detection of misleading information is critical for understanding the persuasive allure of such messages and their power to influence beliefs and behaviour (Pennycook & Rand, 2020). Indeed, such questions have spurred recent research to examine potential mechanisms underlying the transmission and reception of bullshit, finding some cognitive similarities between those who transmit bullshit (i.e., bullshitters) and those who are more receptive to its allure (e.g., Littrell et al., 2020; Pennycook et al., 2015).

Common wisdom suggests that people who frequently mislead others are less likely to be misled themselves, a notion often expressed as, “you can’t bullshit a bullshitter.” This idea finds at least some support in past research showing that people who self-report engaging more frequently in lying (i.e., deliberately convincing someone of a falsehood) also self-report being significantly better than average at detecting lies from others (Zvi & Elaad, 2018). Additionally, some studies have found that those who produce more convincing lies are also actually better at detecting lies (Wright et al., 2012, 2013), though other more recent studies suggest this may not be the case (e.g., Hudson et al., 2020). However, as Frankfurt (2005) and others have pointed out, even though bullshitting is misleading by its very nature, it is distinct from outright deception in that it “falls just short” of lying (e.g., Mears, 2002; Meibauer, 2016). Indeed, recent research has suggested that bullshitting and lying, while clearly related, are psychologically distinguishable constructs (Littrell et al., 2020). For example, liars show a stronger negative association with self-regard and a stronger positive association with lie acceptability than bullshitters (Littrell et al., 2020). Additionally, *persuasive bullshitting* (i.e., bullshitting motivated by a desire to impress or persuade others) has been found to be significantly, negatively related to cognitive ability while

the same has not been found for lying (Littrell et al., 2020; Michels et al., 2020). Given these findings, bullshitters may differ from liars in other meaningful ways, such as their ability to detect the same types of misleading communication that they frequently engage in.

Transmission of bullshit

Two related lines of research have recently emerged investigating individual differences in both the *propensity to produce bullshit* (i.e., bullshitting) and the *propensity to fall for bullshit* (i.e., bullshit receptivity) in a host of situations ranging from social interactions to organizational contexts (Jerrim et al., 2019; McCarthy et al., 2020; Spicer, 2020; Turpin et al., 2019). Here, we define *bullshit*, broadly, as information designed to impress, persuade, and/or otherwise mislead that is often constructed with an indifference for the truth (Carson, 2016; Frankfurt, 2005; Gligorić et al., 2020; Pennycook et al., 2015). Bullshit can range from coherent yet hyperbolic or suspiciously implausible, to jargon-heavy yet obscure or non-sensical, to technically accurate yet misleadingly irrelevant (Carson, 2016; Cohen, 2012; Mears, 2002; Reisch, 2006).

Though the creation of bullshit is intentional, sometimes its spread is not. Indeed, a person might unknowingly or unintentionally transmit bullshit because they mistakenly believe the information to be true. However, as Frankfurt (2006) and others have pointed out, when one engages in *bullshitting*, it is an intentional act by definition (e.g. Cohen, 2012, Mears, 2002; Reisch, 2006). As such, it is largely strategic and utilized to further a goal, such as managing social impressions, increasing status, or influencing opinions. Importantly, though bullshitting is a pervasive aspect of everyday life, only recently have attempts been made to examine its nature empirically. For example, recent work has demonstrated that the extent to which a person intentionally spreads bullshit (i.e., engages in bullshitting) in certain everyday situations can be

estimated using the Bullshitting Frequency Scale (BSF; Littrell et al., 2020). The BSF measures the self-reported frequency with which people strategically engage in producing and transmitting bullshit within various social contexts that is intended to: (1) impress, persuade, or fit in with others by exaggerating, embellishing, or otherwise stretching the truth about one's knowledge, ideas, attitudes, skills, or competence (i.e., *persuasive bullshitting*), and/or; (2) be evasive when responding to inquiries where direct answers might incur negative social costs for oneself or others (i.e., *evasive bullshitting*).

While these two prominent types of bullshitting naturally share significant overlap (i.e., many people encounter and engage in both types of bullshitting in their daily lives), they also differ in a number of important ways including their associations with various cognitive individual differences factors as well as the strategic uses and purposes for which they are employed (Littrell et al., 2020; Mears, 2002; Reisch, 2006). Indeed, recent research suggests that some individuals may be more likely to engage in bullshitting in situations where they believe it will provide them with a social or professional advantage (McCarthy et al., 2020; Spicer, 2020; Turpin et al., 2019). For example, *persuasive bullshitting* is proactively employed to impress or persuade others, often with a loose or casual indifference to the truth of one's statements, such as when an executive makes vacuous, buzzword-heavy embellishments and empty proclamations in an attempt to impress co-workers or influence shareholders (Frankfurt, 2005; Littrell et al., 2020; McCarthy et al., 2020; Spicer, 2020). *Persuasive bullshitting frequency* (as measured by the BSF) has been found to be positively related to performance on tasks thought to reflect "bullshitting behaviour" (Jerrim et al., 2019) such as overclaiming tasks across a range of knowledge domains (i.e., claiming knowledge of non-existent concepts when given the

opportunity) and negatively related to cognitive ability and aspects of analytic thinking (Littrell et al., 2020).

However, not all bullshitting is intended to impress or persuade others. Indeed, in situations where direct answers might result in reputational damage or hurt feelings, a person might reactively engage in *evasive bullshitting*, where the truth is strategically circumnavigated in an attempt to dodge potential social harm (Carson, 2016; Littrell et al., 2020; Meibauer, 2016), such as a politician responding to journalists with evasive, non-relevant truths or strategic ambiguity when asked questions where direct responses might reveal impropriety and/or cost votes (e.g. Cillizza, 2019). Moreover, *evasive bullshitting frequency* has been found to be negatively related to overclaiming and positively related to prosocial lying and providing prosocially evasive (rather than truthful and direct) responses on social decision-making tasks (Littrell et al., 2020).

Receptivity to bullshit

A separate but overlapping line of research exists investigating factors related to *bullshit receptivity*, which refers to the propensity to ascribe inflated judgments of profoundness, truthfulness, or accuracy to information that is vague, obscure, meaningless, or otherwise misleading (Evans et al., 2020; Pennycook et al., 2015; Pennycook & Rand, 2020). Put more simply, it is the tendency to be more receptive to (i.e., fall for) various types of bullshit. Much of the current empirical work in this area has been based on Frankfurt's (2005) notion that bullshit is a type of communication meant to impress and mislead that is often delivered with an indifference for the truth of what one is saying. From this, Pennycook and colleagues (2015) created the Bullshit Receptivity Scale, a collection of statements composed of pseudo-profound

buzzwords that were randomly assembled by an algorithm (thus, indifferent to truth) to be syntactically sound but ultimately meaningless. A higher propensity to rate these types of vacuous statements as profound is negatively associated with cognitive ability and other reflective processes vital for critical thinking and decision-making (Pennycook et al., 2015). Additionally, people higher in *bullshit receptivity* have been found to be more likely to: 1) overclaim their knowledge (Pennycook & Rand, 2019); 2) have a less analytic cognitive style (Evans et al., 2020; Pennycook et al., 2015); 3) detect patterns in patternless images (Walker et al., 2019); 4) give higher profundity ratings to abstract art with randomly-generated names (Turpin et al., 2019), and; 5) endorse various conspiracy theories, such as those related to COVID-19 (Pennycook et al., 2020).

Importantly, researchers have found that people with higher bullshit receptivity are more willing to share pseudo-profound bullshit with others (Čavojová et al., 2018) and more likely to believe and share “fake news” headlines on social media (Pennycook & Rand, 2020). However, it is currently unclear to what extent it is possible that misleading information is transmitted by bullshitters *intentionally* in some instances yet transmitted *unintentionally* in others. If it is indeed the case that bullshitters can themselves be duped by bullshit, this would have important (and potentially nullifying) implications for the utility and effectiveness of bullshitting as a rhetorical persuasion strategy. Although it is arguable whether simply sharing misleading information (that one believes to be true) can be considered a form of “bullshitting,” the fact that evidence exists of a positive relation between belief in and transmission of pseudo-profound bullshit and fake news suggests that some people can be unwitting purveyors of bullshit, a scenario that may ironically extend to bullshitters. However, despite the putative theoretical and correlational overlap between *bullshitting frequency* and *bullshit receptivity*, and their roles in the transmission and reception of

misleading messages, no studies to date have investigated possible associations between these two constructs.

Present investigation

Given that both bullshit receptivity and bullshitting frequency are negatively related to cognitive ability and aspects of analytic thinking style (and that bullshit receptivity is associated with increased sharing of bullshit on social media), it could be the case that those who frequently engage in bullshitting may be more likely to fall for bullshit. However, as noted earlier, separate research suggests that individuals who frequently engage in deception may be better at detecting it, thus leading to the possibility that frequent bullshitters may be less likely to fall for bullshit. Therefore, here we report three studies focused on examining the associations between bullshitting frequency and bullshit receptivity. In Studies 1 and 2, we investigate the correlational and predictive associations among the self-reported propensity to engage in bullshitting, scores on various measures of bullshit receptivity, and performance on measures of a number of cognitive and metacognitive variables. We follow up these correlational results experimentally in Study 3 by investigating potential mechanisms underlying this association.

Data files for all studies are available here: <https://osf.io/chpvm/>.

Study 1

In Study 1, we examine the extent to which *bullshitting frequency* (BSF) is associated with three different types of bullshit: pseudo-profound bullshit, scientific bullshit, and fake news headlines.¹ Each bullshit task also includes a measure of receptivity to contextually-relevant

¹ It should be noted that Study 1 was conducted after Studies 2 and 3. However, we present it here first, as we feel that this presentation order provides helpful theoretical context to the reader for Studies 2 and 3.

nonbullshit (e.g., intentionally profound statements, real scientific information, real news headlines), which allows us to examine the extent to which more frequent bullshitters are receptive to each type of bullshit information while controlling for their receptivity to intentionally profound/scientific/real information. This provides some surface-level insight into their ability to distinguish bullshit from non-bullshit (i.e., their *bullshit sensitivity*).

Method

Participants

We recruited 261 adult participants from the United States and Canada from Amazon's Mechanical Turk participant pool using the crowdsourcing platform, Cloudresearch (Litman, Robinson, & Abberbock, 2016). To meet our goal of achieving at least .80 power to detect an effect of $r = .20$ at $\alpha = .05$, an a priori power analysis indicated that we would need a sample of 191, which our sample exceeded (*g*power*; Faul, Erdfelder, Buchner, & Lang, 2009). Only those who had completed a minimum of 500 surveys and had at least a 97% MTurk HIT approval rating were eligible to participate. Data for 8 participants were removed for being identified by reCAPTCHA v3 bot detection protocols as being potential "bots" while another 34 were removed for failing attention checks, leaving us with data for 219 participants to consider in the final analyses (127 male, 91 female, 1 prefer not to answer, $M_{\text{age}} = 37.94$, $SD_{\text{age}} = 11.44$, Bachelor's degree or higher = 62%). Participants were paid \$2.00 USD for the roughly 15 minute study.

Procedure

After reading an informed consent form, those who agreed to participate answered three demographic questions (i.e., age, biological sex, and level of education). Next, participants completed the remainder of the survey which included the following measures presented in random order (copies of all scale items are listed in the supplementary materials):

Materials

Bullshitting Frequency

To assess the self-reported frequency with which a person utilizes two types of bullshitting in various contexts, we used the Bullshitting Frequency Scale (BSF) from Study 3 of Littrell, Risko, and Fugelsang (2020). Using a 5-point scale ranging from “Never” to “A lot / All the time,” participants rated 12 items by indicating how often they typically engage in bullshitting when confronted with a range of everyday social situations. The BSF comprises two subscales measuring two distinct types of bullshitting; persuasive and evasive. *Persuasive bullshitting* includes attempts to impress, persuade, or fit in with others by exaggerating one’s knowledge, ideas, attitudes, skills, or competence and is measured using items such as, “When I want to contribute to a conversation or discussion even though I’m not well-informed on the topic.” *Evasive bullshitting* is employed to evade/avoid responding to inquiries or situations in which direct answers might incur negative social costs and is measured using items such as, “When being fully honest would be harmful or embarrassing to me or someone else.” Higher scores for each subscale indicate a greater frequency of engaging in that type of bullshitting in certain social contexts and, though it is a self-report measure, these scores have been found to be predictive of performance on tasks involving overclaiming of one’s knowledge and social

decision-making (Littrell et al., 2020). When originally validated, the BSF demonstrated strong reliability for the *persuasive* ($\alpha = .92$) and *evasive* ($\alpha = .82$) subscales.

Pseudo-profound Bullshit Receptivity

In order to assess receptivity to pseudo-profound statements, participants completed the Bullshit Receptivity Scale (BSR; Pennycook et al., 2015) which asks them to evaluate, on a 5point scale (from “not at all profound” to “very profound”), the profundity of 10 randomly generated, yet grammatically correct, sentences that were constructed from abstract pseudoprofound buzzwords (e.g., “We are in the midst of a high-frequency blossoming of interconnectedness that will give us access to the quantum soup itself”). Additionally, participants rated 10 items that represent intentionally profound/motivational quotes (e.g., “A river cuts through a rock, not because of its power but its persistence”).

Scientific Bullshit Receptivity

To measure receptivity to pseudoscientific information, participants completed the Scientific Bullshit Receptivity Scale (SBSR; Evans et al., 2020) where they evaluated the truthfulness (1 = “not at all truthful” to 5 = “very truthful”) of 10 randomly generated, yet grammatically correct, sentences constructed from abstract scientific buzzwords (e.g., “The entropy of an integral approaches constructive interference as its buoyancy approaches endothermal constant of quantum ground states”). Participants also rated 10 statements that convey actual scientific truths (e.g., “In a natural thermodynamic process, the sum of the entropies of the interacting thermodynamic systems increases.”).

Fake News Receptivity

Following procedures from Pennycook and Rand (2020), we presented participants with

10 politically neutral news headlines in picture form as they would appear when posted on social media. Five of the headlines were factually accurate (real news) and five were completely untrue (fake news). Fake news stories were taken from a list of the most popular recent fake news items debunked by Snopes.com. For each headline, participants were asked “To the best of your knowledge, how accurate is the claim in the above headline?” which they indicated on a 4-point scale from “not at all accurate” to “very accurate.” All news headline stimuli can be found in the supplementary materials.

Results

Descriptive statistics and bivariate correlations with persuasive and evasive bullshitting frequency are listed in Table 1 (intercorrelations for all variables can be found in the supplementary materials). Multiple linear regression models were created to examine the extent to which our variables could predict each type of bullshit receptivity (Table 2). We focus first on the correlations and then on the linear regression models.

Correlations

At the bivariate level, *persuasive bullshitting* frequency (BSFp) was significantly and positively related to pseudo-profound bullshit receptivity (BSR), $r(217) = .33, p < .01$, scientific bullshit receptivity, $r(217) = .26, p < .01$, and accuracy ratings of fake news headlines, $r(217) = .36, p < .01$. *Evasive bullshitting* (BSFe) scores were not significantly related to any of our bullshit receptivity measures.

To examine the ways in which persuasive and evasive bullshitting might be differentially related to the other variables, we next calculated partial correlations for each bullshitting type controlling for the other. Given that many people engage in both types of bullshitting (Littrell et

al., 2020), partial correlations allow for better insight into the associations more common to individuals who primarily engage more often in one type of bullshitting over the other. In terms of associations with the bullshit measures, *persuasive bullshitting* was again significantly and positively related to pseudo-profound bullshit receptivity (BSR), $r(216) = .34, p < .01$, scientific bullshit receptivity, $r(216) = .29, p < .01$, and accuracy ratings of fake news headlines, $r(216) = .39, p < .01$. In contrast, the only bullshit measure that *evasive bullshitting* was significantly related to was fake news headlines (negatively), $r(216) = -.18, p = .009$.

Table 1

Descriptive and correlational data for BSF with each type of bullshit receptivity

subscales		<i>M</i>	<i>SD</i>	α	<i>Partials</i>			
					<i>Bivariate</i>		<i>controlling for BSF</i>	
					BSFp	BSFe	BSFp ^a	BSFe ^b
1	Persuasive bullshitting (BSFp)	2.51	0.88	.92	-	-	-	-
2	Evasive bullshitting (BSFe)	3.03	0.79	.81	.53**	-	-	-
3	Pseudo-profound bullshit	27.61	10.42	.94	.33**	.08	.34**	-.13
4	Profound motivational quotes	34.23	7.68	.83	.30**	.18**	.24**	.03
5	Scientific bullshit	31.56	6.53	.82	.26**	.04	.29**	-.12
6	Real scientific statements	34.05	5.73	.73	.18**	.11	.14*	.01
7	Fake news headlines	8.37	2.89	.70	.36**	.05	.39**	-.18*
8	Real news headlines	14.37	2.72	.62	-.02	.11	-.09	.14*

Note: $N = 219$. BSF = Bullshitting Frequency Scale; BSFp^a = Persuasive bullshitting, controlling for evasive; BSFe^b = Evasive bullshitting, controlling for persuasive. ** $p < .01$; * $p < .05$

Linear regressions

To test our main question of whether propensity to engage in bullshitting predicts receptivity to various types of bullshit, we created three multiple linear regression models, each predicting one of the three bullshit receptivity scores (i.e., pseudo-profound, scientific, and fake news) from bullshitting frequency scores (persuasive and evasive). As a third predictor in each model, we also included the relevant non-bullshit measure for each of the bullshit receptivity tasks. Doing so allows us to use multiple linear regression to utilize our bullshit receptivity variable (i.e., the extent to which one is generally receptive to bullshit) as an index of *bullshit insensitivity* (i.e., one's inability to distinguish bullshit from non-bullshit) by controlling for one's general receptivity to contextually similar, non-misleading information. This method offers an arguably less biased measure of insensitivity than calculated difference scores and has been recommended in past research as a suitable alternative when assessing sensitivity/insensitivity for continuous variables using multiple linear regression analysis (Edwards, 1994, 1995; Cafri et al., 2010; Peter et al., 1993; Vickers & Altman, 2001; for a more fulsome discussion, see Belmi et al., 2020). However, for clarity and ease of interpretation, each outcome measure will be referred to as an index of *receptivity* for each specific type of misleading information.²

Pseudo-profound bullshit receptivity (BSR) was significantly and positively predicted by persuasive bullshitting frequency (BSFp), $\beta = .28$, $p < .01$, 95% CI [.14, .41] and profound/motivational quote receptivity (MQR), $\beta = .47$, $p < .01$, 95% CI [.35, .59], and significantly and negatively predicted by evasive bullshitting frequency (BSFe), $\beta = -.16$, $p = .02$,

² We also conducted separate analyses using difference scores (e.g., subtracting BSR from MQR) as a dependent variable in order to confirm our results and found the patterns of associations to be identical to those reported here.

95% CI [-.29, -.03]. Likewise, *scientific bullshit receptivity* was significantly and positively predicted by persuasive bullshitting frequency (BSFp), $\beta = .23$, $p < .01$, 95% CI [.12, .35] and

receptivity to real science information (MQR), $\beta = .62$, $p < .01$, 95% CI [.52, .72], and significantly and negatively predicted by evasive bullshitting frequency (BSFe), $\beta = -.15$, $p = .01$, 95% CI [-.27, -.03]. Finally, *fake news receptivity* was significantly and positively predicted by persuasive bullshitting frequency (BSFp), $\beta = .47$, $p < .01$, 95% CI [.32, .61] and significantly and negatively predicted by evasive bullshitting frequency (BSFe), $\beta = -.21$, $p = .006$, 95% CI [.35, -.06]. However, it was not significantly predicted by accuracy judgements (i.e., receptivity) of real news headlines (BSFp), $\beta = .07$, $p = .24$, 95% CI [-.05, .20].

Table 2

Study 1 multiple linear regressions for bullshitting frequency predicting receptivity to each bullshit type

	Bullshit receptivity		
	BSR	SBSR	FNR
Persuasive bullshitting (BSFp)	.28**	.23**	.47**
Evasive bullshitting (BSFe)	-.16*	-.15*	-.21**
Profoundness receptivity	.47**		
Real science truthfulness		.62**	
Real news headline accuracy			.07
Adjusted R^2	.32	.45	.15
F	34.47**	60.83**	13.56**

Note: $N = 219$. Standardized beta coefficients listed. BSR = Bullshit Receptivity Scale; SBSR = Scientific bullshit receptivity scale; FN = Fake news headline receptivity

** $p < .01$; * $p < .05$

Discussion

In Study 1, we examined the extent to which individual differences in the propensity to engage in two types of bullshitting (i.e., bullshitting frequency) are related to the propensity to fall for three different types of bullshit (i.e., bullshit receptivity). Correlational results indicated that *persuasive bullshitting* (but not evasive) was positively associated with receptivity to pseudo-profound bullshit, scientific bullshit, and fake news headlines. Crucially, our linear regression models revealed that the frequency with which a person engages in *persuasive* and *evasive bullshitting* significantly predicts receptivity to each type of bullshit, even when controlling for receptivity to contextually-relevant non-bullshit information. It is also noteworthy that the direction of the associations with bullshit receptivity between both types of bullshitting frequency diverged in opposing directions. This provides some evidence that the two types of bullshitting may rely on different cognitive processes, which has been suggested in prior research (Littrell et al., 2020). Overall, these findings provide evidence that more frequent persuasive bullshitters are more susceptible to falling for various types of misinformation. This has important implications regarding the ways in which some types of misleading information are transmitted and received, in that some people who spread misleading information intentionally (e.g., bullshitters) may also be susceptible to spreading – and falling for – it unintentionally, as they have difficulty discerning fact from fiction.

Study 2

In Study 2, we examine more deeply the extent to which bullshitting frequency is associated with receptivity to misleading information (i.e., bullshit). To that end, we limit our examination to one type of bullshit (i.e., pseudo-profound) as measured by the BSR. Importantly, the BSR has been found in multiple studies to be correlated with higher endorsement of fake

news headlines, receptivity to pseudo-scientific information, receptivity to empty and misleading political statements, beliefs in COVID-19 conspiracies, and a less analytic thinking style when evaluating information and problem-solving (Evans et al., 2020; Gligorić et al., 2020; Pennycook et al., 2020; Pennycook & Rand, 2020). Given this, we felt that receptivity to pseudo-profound bullshit is a good proxy for receptivity to a wide range of various types of misleading information and epistemically suspect beliefs (i.e., general bullshit).

We also include a number of additional measures to allow us to explore the nature of the relation between bullshitting frequency and bullshit receptivity, focusing on three general classes of cognitive predictors: 1) cognitive ability; 2) factors related to engagement in various facets of cognitive reflection, and; 3) subjective and objective measures of metacognition. The first two classes of variables were selected given their established relation to bullshitting frequency and bullshit receptivity (Littrell et al., 2020; Pennycook et al., 2015). The third was selected based on findings that bullshitting is related to overconfidence (Jerrim et al., 2019). Thus, these variables represent potential mediators of the putative relation between bullshitting frequency and pseudoprofound bullshit receptivity.

Method

Participants

We recruited 210 adult participants from the United States and Canada from Amazon's Mechanical Turk participant pool using the crowdsourcing platform, Cloudresearch (Litman, Robinson, & Abberbock, 2016). This was based on our goal of achieving at least .80 power to detect an effect of $r = .20$ at $\alpha = .05$ (g*power; Faul, Erdfelder, Buchner, & Lang, 2009). Only those who had completed a minimum of 100 surveys and had at least a 95% MTurk HIT approval rating were eligible to participate. Data for three participants was removed for failing attention

checks, leaving us with data for 207 participants to consider in the final analyses (137 male, 69 female, 1 intersex, $M_{\text{age}} = 36.75$, $SD_{\text{age}} = 11.18$, Bachelor's degree or higher = 60.9%).

Participants were paid \$3.50 USD for the roughly 25 minute survey.

Procedure and Materials

Informed consent and online survey presentation procedures were the same as Study 1. Participants again completed the BSR as well as a version of the BSF from Study 1 of Littrell et al. (2020)³. Additionally, participants completed the following measures presented in random order:

Cognitive ability

To assess participants' ability to understand and carry out basic mathematical operations, a 10-item version of the General Risk and Numeracy Scale was administered (Lipkus, Samsa, & Rimer, 2001). Verbal intelligence was assessed with a 10-item version of the "Wordsum" vocabulary test (Thorndike, 1942; Malhotra, Krosnick, & Haertel, 2007). Scores on both tests were combined to calculate a *mean cognitive ability* score.

We also collected confidence ratings for each cognitive ability item, using a sliding scale from 0-100, from which we calculated a *cognitive ability bias* score for each participant (i.e., an objective measure of intellectual overconfidence). Mean cognitive ability scores were converted to percentages and then subtracted from the average confidence score for the cognitive ability items to give an index of bias. Scores above zero indicate intellectual *overconfidence* while scores falling below zero indicate intellectual *underconfidence*.

Cognitive Reflection

³ See supplementary materials.

To assess participants' ability to reflectively override conflict during problem-solving, participants completed a 10-item version of the Cognitive Reflection Test (CRT-10). The CRT-10 consists of 10 "brain teasers"; three from Frederick's (2005) original CRT, three items added by Primi et al (2016), and four taken from Thomson and Oppenheimer (2016). Additionally, participants' self-reported engagement in cognitive reflection was measured using twelve items from Grant, Franklin, and Langford's (2002) Self-Reflection and Insight Scale (SRIS), which represent a person's need and propensity to reflect on and evaluate their thoughts, feelings, and behaviours. Participants rated themselves on items such as, "I frequently take time to reflect on my thoughts" using a 5-point Likert scale from "Strongly disagree" to "Strongly agree." Grant et al., (2002) reported excellent reliability for the reflection subscale in the original validation ($\alpha = .91$).

Metacognition (self-reported)

We assessed the degree to which participants report clearly understanding their own thoughts, feelings, and behaviours (i.e., their self-reported *metacognitive insight*) using 8-items from the Self-Reflection and Insight Scale (Grant et al., 2002). Participants rated items such as, "Thinking about my thoughts makes me more confused," using a 5-point Likert scale from "strongly disagree" to "strongly agree." The insight subscale has shown excellent reliability in past research ($\alpha = .90$; Littrell, Fugelsang, & Risko, 2020).

Self-reported intellectual overconfidence was assessed using the *lack of intellectual overconfidence* subscale of Krumrei-Mancuso's and Rouse's (2016) Comprehensive Intellectual Humility Scale. Participants rated themselves on a 5-point Likert scale using items such as,

“When I am really confident in a belief, there is very little chance that belief is wrong.” The original validation study reported an acceptable average internal reliability for this subscale ($\alpha = .72$; Krumrei-Mancuso & Rouse, 2016). In order to capture self-reported *intellectual overconfidence*, we reverse-scored the scale.

Results

Correlations

Correlations and descriptive statistics for key variables can be found in Table 3 (intercorrelations for all variables can be found in the supplementary materials). Both *persuasive bullshitting* (BSFp), $r(205) = .39, p < .01$, and *evasive bullshitting*, $r(205) = .22, p < .01$, were significantly related to pseudo-profound bullshit receptivity (BSR). *Persuasive bullshitting* was also positively related to self-reported intellectual overconfidence, $r(205) = .30, p < .01$, and calculated intellectual overconfidence (i.e., bias), $r(205) = .25, p < .01$, and negatively related to cognitive ability, $r(205) = -.25, p < .01$, CRT scores, $r(205) = -.20, p < .01$, and insight, $r(205) = -.42, p < .01$.

Receptivity to pseudo-profound bullshit (BSR) was significantly and positively related to self-reported overconfidence, $r(205) = .24, p < .01$, and calculated overconfidence, $r(205) = .27, p < .01$, and significantly and negatively related to cognitive ability, $r(205) = -.42, p < .01$, CRT, $r(205) = -.44, p < .01$, and insight, $r(205) = -.21, p < .01$.

As with Study 1, to better understand the associations for individuals who more often engage in one type of bullshitting over the other, we next calculated partial correlations for each bullshitting type (BSFp and BSFe) controlling for the other. *Persuasive bullshitting* was positively related to bullshit receptivity, $r(204) = .34, p < .01$, self-reported intellectual overconfidence, $r(204) = .32, p < .01$, and calculated overconfidence, $r(204) = .26, p < .01$, and

negatively related to cognitive ability, $r(204) = -.31, p < .01$, CRT scores, $r(204) = -.19, p < .01$, and insight, $r(204) = -.32, p < .01$. In contrast, *evasive bullshitting* was significantly related only to cognitive ability (positively), $r(204) = .20, p < .01$, and self-reported overconfidence (negatively), $r(204) = -.17, p = .02$.

Table 3

Descriptive and correlational data for BSF with all study variables

	<i>M</i>	<i>SD</i>	α	<i>Bivariate</i>			<i>Partial</i>	
				BSFp	BSFe	BSR	BSFp ^a	BSFe ^b
1 Persuasive bullshitting (BSFp)	2.49	0.82	.92	-			-	-
2 Evasive bullshitting (BSFe)	2.79	0.83	.82	.74**	-		-	-
3 Bullshit receptivity (BSR)	25.64	9.72	.92	.39**	.22**	-	.34**	-.10
4 Profoundness receptivity (MQR)	34.19	7.33	.86	.21**	.22**	.48**	.07	.10
5 Cognitive ability	8.10	1.72	.84	-.25**	-.05	-.42**	-.31**	.20**
6 Cognitive reflection test (CRT)	7.15	2.71	.84	-.20**	-.10	-.44**	-.19**	.07
7 Self-reported reflection	3.68	0.84	.94	-.06	.01	.06	-.10	.08
8 Insight	3.77	0.77	.87	-.42**	-.29**	-.21**	-.32**	.03
9 Overconfidence (self-reported)	2.78	0.85	.85	.30**	.12	.24**	.32**	-.17*
10 Overconfidence (calculated)	3.37	16.45	-	.25**	.10	.27**	.26**	-.12

Note: $N = 207$. BSF = Bullshitting Frequency Scale; BSFpa = Persuasive bullshitting, controlling for evasive; BSFeb = Evasive bullshitting, controlling for persuasive.

** $p < .01$; * $p < .05$

Linear regression

We next created a multiple linear regression model predicting *bullshit receptivity* from bullshitting frequency scores (persuasive and evasive) while entering the remaining variables as covariates. Predictors were entered in three steps. In Step 1, we entered our two bullshitting frequency variables and receptivity to motivational quotes. In Step 2, we entered cognitive ability. Finally, in Step 3, we entered our cognitive reflection and metacognition variables (i.e., CRT, insight, and our two overconfidence variables). As self-reported reflection was not significantly associated with bullshit receptivity at the bivariate level, it was excluded from our regression model. Our discussion will focus only on the final, overall model (i.e., Step 3),

however, standardized beta coefficients and model fit statistics for each step of the regression are listed in Table 4.

In the final model, bullshit receptivity was significantly and positively predicted by persuasive bullshitting frequency (BSFp), $\beta = .32$, $p < .01$, 95% CI [.14, .50], and receptivity to motivational quotes (MQR), $\beta = .37$, $p < .01$, 95% CI [.25, .48], and significantly and negatively predicted by cognitive ability $\beta = -.32$, $p = .004$, 95% CI [-.53, -.10]. Additionally, though scores on the CRT, $\beta = -.15$, $p = .063$, 95% CI [-.30, -.008], and the calculated intellectual overconfidence measure, $\beta = -.17$, $p = .053$, 95% CI [-.35, -.002], were negative predictors of bullshit receptivity, both fell just short of statistical significance.

Table 4

Study 2 multiple linear regressions for all study variables predicting receptivity to bullshit

	Bullshit Receptivity		
	Step 1	Step 2	Step 3
Persuasive bullshitting (BSFp)	.44**	.31**	.32**
Evasive bullshitting (BSFe)	-.20*	-.11	-.10
Profoundness receptivity (MQR)	.43**	.39**	.37**
Cognitive ability		-.29**	-.32**
Cognitive Reflection Test			-.15
Insight			.05
Overconfidence (self-reported)			.06
Overconfidence (calculated)			-.17
Adjusted R^2	.32	.39	.41
F	33.30**	34.00**	18.80**

Note: $N = 207$. Standardized beta coefficients listed.

** $p < .01$; * $p = .05$

Discussion

Study 2 delved deeper into the association between the propensity to engage in bullshitting and the propensity to fall for bullshit. As with Study 1, partial correlations revealed that persuasive bullshitting (positive) and evasive bullshitting (negative) were differentially related to bullshit receptivity. Indeed, the two types of bullshitting (when controlling for the other) diverged in opposite directions on a number of our cognitive and metacognitive variables. For instance, persuasive bullshitting (controlling for evasive) was *positively* related to bullshit receptivity and our overconfidence measures, while evasive (controlling for persuasive) was *negatively* related to these factors. Additionally, persuasive bullshitting was *negatively* related to cognitive ability, CRT, and insight whereas evasive bullshitting was *positively* related to cognitive ability and unrelated to CRT scores or insight. This provides more evidence that different cognitive profiles may underlie the proclivities of some people to primarily engage in one type of bullshitting over the other. That is, persuasive bullshitting may rely on less engagement in analytic thinking processes compared to evasive bullshitting.

Importantly, our linear regression model supports Study 1's findings that the frequency with which a person engages in *persuasive bullshitting* positively predicts bullshit receptivity, even when potential mediators of such a relation (e.g., evasive bullshitting, overall profoundness receptivity, metacognitive ability, and cognitive ability) are taken into account. Thus, the relation between persuasive bullshitting and bullshit receptivity does not appear to be explained by a tendency to see profoundness everywhere, the propensity and/or capability to reflect, perceived clarity of thought, overconfidence, or cognitive ability. This was true despite the fact that both persuasive bullshitting frequency and bullshit receptivity were related to these constructs in a theoretically consistent manner at the bivariate level.

Additionally, *evasive bullshitting* frequency negatively predicted bullshit receptivity in Step 1 of our linear regression model, even after controlling for overall profoundness receptivity (i.e., profound motivational quotes). That is, while people who primarily engage in *persuasive bullshitting* were *more likely* to fall for bullshit, people who primarily engage in *evasive bullshitting* appeared *less likely* to fall for bullshit. Notably, controlling for cognitive ability in Step 2 reduced this relation to non-significant. This suggests that the negative relation between evasive bullshitting and bullshit receptivity may be at least partially due to an individual's intelligence. However, it is important to note that many people self-report a tendency to engage in both types of bullshitting at fairly equal frequencies across various contexts, therefore other individual differences that we have not accounted for may play important roles in the extent to which these individuals are receptive to misleading information.

Overall, these results further support the idea that bullshitting frequency predicts bullshit receptivity (insensitivity) and the type of bullshitting determines the direction of this association. Furthermore, certain facets of metacognitive processes and cognitive ability are related to both the propensity to produce and the propensity to be receptive to bullshit, which is consistent with previous work (Littrell et al., 2020; Pennycook et al., 2015). But, these processes (at least as measured here), do not appear to underlie (completely) the relation between bullshitting frequency and bullshit receptivity, particularly in the case of persuasive bullshitting. That is, controlling for these variables did not eliminate the positive relation between persuasive bullshitting and bullshit receptivity (though it was slightly diminished) but there was some evidence that doing so eliminated the negative relation between evasive bullshitting and bullshit receptivity.

Study 3

Studies 1 and 2 have established a consistent, robust positive association between persuasive bullshitting frequency and bullshit receptivity, but the mechanism underlying this relation remains unclear. In Study 3, we test one potential mechanism for this association; that is, whether it is the product of a relative insensitivity in higher frequency persuasive bullshitters to the differences between statements that “sound profound” and those that actually “are profound.” To do this, we devised two new sets of BSR instructions. One set asked participants to rate BSR items based on *how profound they sound*, ignoring how profound they believe the items actually are. The other set of instructions reversed this, asking participants to rate the items based on *how profound they actually are*, ignoring how profound they subjectively sound. If the positive relation between persuasive bullshitting and pseudo-profound bullshit receptivity is based (to some extent) on individuals high in persuasive bullshitting being insensitive (relative to those low in persuasive bullshitting) to the distinction between statements “sounding profound” and actually “being profound,” then there should be an interaction with instruction and persuasive bullshitting. The form of this interaction should reflect the instruction having a minimal effect on individuals high in persuasive bullshitting relative to those low in persuasive bullshitting.⁴

Participants

To ensure that we achieved at least .80 power to detect an effect of $r = .20$ at $\alpha = .05$ in both conditions, we recruited 454 adult participants from the United States and Canada from

⁴ Prior to Study 3, we conducted a pre-registered study in which we tested a more subtle instruction manipulation. These results showed that a lower propensity to engage in cognitive reflection does not explain the positive predictive association between persuasive bullshitting frequency and bullshit receptivity. Given the provided word count restrictions, and in an effort to reduce redundancy, we chose to report those results in full in the supplementary materials as Study 2b.

Amazon's Mechanical Turk using Cloudresearch (Litman, Robinson, & Abberbock, 2016). Data for 54 participants were eliminated for failing attention checks or being identified as potential "bots" by reCAPTCHA v3 bot detection protocols, leaving us with a sample of 400 to consider in the final analyses. Participants were paid \$0.75 USD for the roughly 10 minute study.

Procedure

Our hypotheses and methods were preregistered on OSF and are available at <https://osf.io/3k6tn>. Recruitment and survey administration procedures were the same as those of the previous studies.

Materials

The materials used to measure bullshitting frequency and cognitive ability were identical to Study 2. Participants were randomly assigned to one of two BSR instruction conditions. The instructions for Group 1 were: "*We are interested in what makes items profound. We would like you to rate the following items with respect to how profound they sound. Please ignore how profound you think each statement truly is (i.e., how profound an item sounds might not be related to how profound that statement is). The definition of profound is 'showing great knowledge or insight; to be taken as deeply meaningful; of great and broadly inclusive significance.'*"

Group 2 received these instructions: “*We are interested in what makes items profound. We would like you to rate the items below with respect to how profound you think each statement truly is. Please ignore how profound each statement sounds (i.e., how profound an item sounds might not be related to how profound that statement is). The definition of profound is ‘showing great knowledge or insight; to be taken as deeply meaningful; of great and broadly inclusive significance.’*” A reminder of the wording of instructions was presented to both groups a second time on a screen by itself, directly after the primary instruction screen, to decrease the chances that participants skimmed over the instructions without reading them.

Results

We focus our results and discussion on the linear regression analyses of the experimental manipulation. However, descriptive statistics and intercorrelations for all variables can be found in the supplementary materials. Following previous procedures, we created a multiple linear regression model predicting overall BSR scores (i.e., bullshit insensitivity) from persuasive bullshitting (BSFp), evasive bullshitting (BSFe), profound/motivational quote receptivity (MQR), and cognitive ability (Table 5). We included a variable for BSR instruction condition and interaction terms for both of our bullshitting frequency variables (see Figure 1). All predictor variables (excluding condition) were mean-centred.

Neither persuasive, $\beta = .09$, $p = .26$, 95% CI [-.07, .25] nor evasive, $\beta = -.14$, $p = .09$, 95% CI [-.30, .02], bullshitting significantly predicted bullshit receptivity on their own. For the interaction variables, the PersuasiveBS*condition interaction significantly and positively predicted bullshit insensitivity, $\beta = .18$, $p = .03$, 95% CI [.02, .35], though the EvasiveBS*condition interaction was not significant. Following up on this finding, we created regression models for each instruction condition (Figure 1). This revealed that in the “is

profound” condition, both persuasive, $\beta = .37$, $p < .001$, 95% CI [.21, .53], and evasive, $\beta = -.31$, $p < .001$, 95% CI [-.47, -.15], bullshitting significantly predicted bullshit receptivity (controlling for MQR and cognitive ability), matching the pattern found in Studies 1 and 2 for the standard instructions. However, neither persuasive, $\beta = .12$, $p = .19$, 95% CI [-.06, .30], nor evasive, $\beta = .15$, $p = .08$, 95% CI [-.33, .02], bullshitting were significant predictors of bullshit receptivity in the “sounds profound” condition, though their associations with BSR trended in expected directions. A closer inspection of the slopes for each condition (Figure 1) reveals that, when controlling for the other variables, individuals scoring low in persuasive bullshitting rated pseudo-profound items lower in the “is profound” condition compared to the “sounds profound” condition, while individuals scoring high in persuasive bullshitting gave higher profoundness ratings to pseudo-profound items in both conditions. It should also be noted that, individuals higher in evasive bullshitting rated items in the “is profound” condition lower than in the “sounds profound” condition, though this interaction failed to reach statistical significance.

Table 5

Study 3 multiple linear regressions for all study variables (mean-centred) predicting bullshit receptivity (BSR score)

	<i>b</i>	<i>SE</i>	β	95% CI Lower Upper	
<i>(Constant)</i>	29.44**	0.55			
Persuasive bullshitting	1.02	0.90	.09	-.07	.25
Evasive bullshitting	-1.57	0.92	-.14	-.30	.02
Profoundness receptivity (MQR)	.52**	0.05	.41**	.32	.49
Cognitive ability	-1.74**	0.29	-.25**	-.33	-.17
Condition	-3.19**	0.78	-.17**	-.25	-.09
PersuasiveBS*condition	2.91*	1.30	.19*	.02	.35
EvasiveBS*condition	-1.78	1.31	-.11	-.28	.05
Adjusted R^2	.33				
<i>F</i>	28.98**				

Note: $N = 400$. CI = confidence intervals for standardized betas; $**p < .01$; $*p < .05$




Figure 1. Plots of interaction effects for persuasive and evasive bullshitting predicting bullshit receptivity (BSR), controlling for all other variables. Dark solid lines represent “Sounds profound” instruction condition. Light dashed lines represent “Is profound.” instruction condition.

Discussion

Our goal for Study 3 was to examine whether the positive association between persuasive bullshitting and bullshit receptivity could be explained by a failure among high persuasive bullshitters to meaningfully distinguish between items that simply “sound profound” and items that arguably “are profound” (or at least generally accepted to be so). Based on the results presented here, it appears that high persuasive bullshitters do struggle in making this distinction while high evasive bullshitters do not, even when cognitive ability is taken into account. Importantly, our linear regression analyses showed that the interaction between *persuasive bullshitting* and condition with bullshit receptivity was significant. Specifically, individuals scoring lower in persuasive bullshitting gave lower profoundness ratings in the “is profound” condition while those higher in persuasive bullshitting gave higher profoundness ratings (compared to low bullshitters) to these statements. Additionally, high persuasive bullshitters rated the statements in both conditions (i.e., “sounds” and “is”) as approximately equally profound. Put another way, high *persuasive bullshitters* appear to interpret/mistake *superficial profoundness* as a signal of *actual profoundness*. Conversely, while low *evasive bullshitters* tended to rate items that “sounded” profound on approximately equal par with those they deemed to actually be profound, high evasive bullshitters were clearly better able to distinguish between “sounding profound” and “being profound.”

General discussion

Across three studies ($N = 826$), we found consistent support for a positive association between persuasive bullshitting frequency and susceptibility to falling for various types of misleading information (e.g. pseudo-profound bullshit, scientific bullshit, and fake news).

Additionally, *evasive* bullshitting was negatively associated with receptivity to these same types of misleading information (though this negative association was non-significant in some instances). Furthermore, the predictive association between persuasive bullshitting and pseudoprofound bullshit receptivity was robust in that it was largely unaffected when controlling for potential cognitive and metacognitive mediators thought to underlie this association. Importantly, we found evidence that people high in persuasive bullshitting appear unable to distinguish *superficial profundness* (i.e., a statement simply “sounding profound”) from *inherent profundness* (i.e., actually “being profound”). In other words, for persuasive bullshitters, if a statement *sounds* profound, to them that indicates that the statement truly *is* profound. In contrast, high evasive bullshitters (compared to persuasive) seem better equipped to make this distinction. In some ways, this appears to somewhat align with research suggesting that individuals more willing to share fake news (in some instances) are also more likely to fall for it (Pennycook & Rand, 2019), but also appears to somewhat conflict with other research suggesting a positive relation between lying and lie detection (Wright et al., 2012; Zvi & Elaad, 2018). These findings support the idea that being more likely to produce bullshit does not necessarily inoculate a person from being more likely to fall for bullshit (i.e., one *can* “bullshit a bullshitter”). In the following, we expand upon these findings and suggest some potentially fruitful directions for future research.

A bullshit blindspot

Implicit within the observations presented here are the somewhat complex interpersonal dynamics involved in how bullshit is produced, transmitted, and received. As Frankfurt (2006) and others have defined it, bullshitting is intentional, deliberate, and strategic (Littrell et al.,

2020; Mears, 2002; Reisch, 2006). For example, a person can massage truthful information in a way that would be, by definition, “bullshitting” if he is doing so to be misleading or misrepresent his own goals (Frankfurt, 2005; Mears, 2002; Reisch, 2006). However, if a bullshitter transmits information in an earnest attempt to convey a true message, yet is unaware the information he is transmitting is actually bullshit, he is not (by definition) engaging in “bullshitting” because there was no intention to mislead or misrepresent by statement or implicature (Frankfurt, 2005; Meibauer, 2016; Webber, 2013).

Consequently, just as a liar might unknowingly “spread lies” (because he believes them to be true), he cannot unintentionally *engage in lying*. Likewise, a bullshitter might unknowingly “spread bullshit” (because he believes it to be true) but cannot unintentionally *engage in bullshitting*. This has important implications regarding the extent to which bullshitters are able to recognize (and possibly prevent) those times when they are unknowingly spreading bullshit. Given the intentional, strategic nature of bullshitting, if a bullshitter unintentionally or unknowingly spreads bullshit at a strategically disadvantageous time (because he or she is unable to detect it), it may nullify both the perceived and actual utility of bullshitting as a rhetorical persuasion strategy for that person in general.

We attempted to address this issue in the present study, at least in part, by testing the “bullshit insensitivity” abilities of two types of self-reported prolific bullshitters with empirical measures of various types of bullshit receptivity. One limitation, though, is that we did not ask participants to assess their own “bullshit detection” abilities, as previous deception research has done (e.g., Zvi & Elaad, 2018). Indeed, given that higher frequency persuasive bullshitters were (somewhat ironically) consistently found to be more receptive to various types of bullshit, and

were simultaneously overconfident in their own intellectual abilities, it could very well be the case that they are largely unaware of their own inability to sufficiently detect when they are being misled. That is, higher frequency persuasive bullshitters may experience unique DunningKruger-like effects related to their own perceived and actual ability to detect misleading information (Pennycook et al., 2017). Put another way, they may have a “bullshit blind spot” akin to that found in other domains (Pronin, Lin, & Ross, 2002). Therefore, it would be informative for future bullshitting research to investigate the extent to which the self-assessed and empirically-measured bullshit detection abilities of persuasive bullshitters align, as well as how bullshit-specific overconfidence might be related to other analytic and metacognitive processes that play important roles in the transmission and detection of various types of misleading information.

Bullshitting frequency, bullshitting quality, and intelligence

Another finding presented here that may seem counterintuitive given past research on deception, is the negative relation between persuasive bullshitting and intelligence. Indeed, past work has asserted that people of higher intelligence should be more adept at strategically misleading others (Handel, 1982). However, research into the deception abilities of prolific liars has thus far not identified a meaningful connection between objective deception ability and intelligence in adults (e.g., Michels et al., 2020). For instance, Wright, Berry, and Bird (2012, 2013) found that people who were able to produce more convincing lies (i.e., better liars) were also better able to detect lies from others but that this ability was not significantly related to intelligence. Conversely, preliminary work investigating the relation between bullshit production and intelligence suggests that people who are more intelligent are able to produce more

convincing bullshit compared to people of lower intelligence and that this “bullshit production ability” may be *unrelated* to bullshit receptivity (Turpin et al., 2020).

However, these previous studies did not measure the *frequency* with which participants self-report engaging in lying/bullshitting, and our investigation did not address bullshitting *quality*. Given the present results as well as previous work (Littrell et al., 2020) showing a negative association between persuasive bullshitting frequency and intelligence, it may be the case that *bigger bullshitters* are not necessarily *better bullshitters*. Indeed, less intelligent people may be more likely to find themselves in situations in which they feel intellectually underprepared yet still desire to leverage attitudes and impressions in their favour. In these situations, they may engage in a higher frequency of persuasive bullshitting but lack the cognitive and intellectual horsepower to produce bullshit that is convincing. Conversely, people of higher intelligence would be more likely to possess the requisite cognitive and intellectual faculties to produce higher quality, more convincing bullshit but may paradoxically engage in such behaviour less frequently, as they would be less likely to experience situations in which they feel intellectually outmatched. Bringing these related lines of research together seems a logical “next step” for future bullshitting research to take.

Persuasive versus evasive bullshitting

Finally, the present results provide more evidence of the cognitive and individual differences between persuasive and evasive bullshitting frequency. As demonstrated here and in previous research (Littrell et al., 2020), *persuasive bullshitting* is negatively related to cognitive ability and analytic thinking and positively related to overclaiming and overconfidence. Conversely, *evasive bullshitting* is positively related to cognitive ability and negatively related to overclaiming and overconfidence. Importantly, the present results show that higher frequency

persuasive bullshitters are more receptive to misleading information while higher frequency *evasive bullshitters are less receptive* to misleading information (i.e., bullshit). Though more research is still needed, the emerging distinctions between persuasive and evasive bullshitting appear to fit with the *promotion focus vs. prevention focus* (or *approach vs. avoidance*) goal pursuit distinctions found within the motivational and self-regulatory literature (e.g., Higgins, 2012). Indeed, the two types of bullshitting serve different strategic purposes and appear to be motivated by different situational and interpersonal factors, therefore a deeper exploration of their differences from a motivational perspective would likely be a fruitful line of future inquiry.

Conclusion

Gaining a better understanding of the differing ways in which various types of misleading information are transmitted and received is becoming increasingly important in the information age (Kristansen & Kaussler, 2018). Indeed, an oft-repeated maxim in popular culture is, “you can’t bullshit a bullshitter.” While folk wisdom may assert that this is true, the present investigation suggests that the reality is a bit more complicated. Our primary aim was to examine the extent to which bullshitting frequency is associated with susceptibility to falling for bullshit. Overall, we found that persuasive bullshitters (but not evasive bullshitters) were more receptive to various types of bullshit and, in the case of pseudo-profound statements, even when controlling for factors related to intelligence and analytic thinking. These results enrich our understanding of the transmission and detection of certain types of misleading information, specifically the associations between the propensity to produce and the tendency to fall for bullshit and will help to inform future research in this growing area of scholarship.

References

- Belmi, P., Neale, M. A., Reiff, D., & Uffe, R. (2020). The social advantage of miscalibrated individuals: The relationship between social class and overconfidence and its implications for class-based inequality. *Journal of Personality and Social Psychology*, *118*(2), 254-282. <http://dx.doi.org/10.1037/pspi0000187>
- Cafri, G., van den Berg, P., & Brannick, M. T. (2010). What Have the Difference Scores Not Been Telling Us? A Critique of the Use of Self—Ideal Discrepancy in the Assessment of Body Image and Evaluation of an Alternative Data-Analytic Framework. *Assessment*, *17*(3), 361–376. <https://doi.org/10.1177/1073191109357028>
- Carson, T. (2016). Frankfurt and Cohen on bullshit, bullshitting, deception, lying, and concern with the truth of what one says. *Pragmatics & Cognition* *23*, 54-68. <https://doi.org/10.1075/pc.23.1.03car>
- Čavojová, V., Secară, E., Jurkovič, M., & Šrol, J. (2019). Reception and willingness to share pseudo-profound bullshit and their relation to other epistemically suspect beliefs and cognitive ability in Slovakia and Romania. *Applied Cognitive Psychology*, *33*(2), 299-311. <https://doi.org/10.1002/acp.3486>
- Cillizza, C (2019, October 1). *Mike Pompeo gave a very evasive answer when asked about the Ukraine call*. CNN. <https://www.cnn.com/2019/10/01/politics/mike-pompeo-ukraine/index.html>
- Cohen, G. (2012). Complete bullshit. In Otsuka M. (Ed.), *Finding Oneself in the Other* (pp. 94-114). Princeton University Press. Retrieved from <http://www.jstor.org/stable/j.cttq956b.9>
- Edwards, J. R. (1994). Regression Analysis as an Alternative to Difference Scores. *Journal of*

- Management*, 20(3), 683–689. <https://doi.org/10.1177/014920639402000311>
- Edwards, J. R. (1995). Alternatives to difference scores as dependent variables in the study of congruence in organizational research. *Organizational Behavior and Human Decision Processes*, 64(3), 307–324. <https://doi.org/10.1006/obhd.1995.1108>
- Evans, A. M., Slegers, W., & Mlakar, Ž. (2020, February 21). Individual differences in receptivity to scientific bullshit. *Judgment and Decision-Making*, 15(3), 401-412. <http://journal.sjdm.org/20/200221/jdm200221.pdf>
- Faul, F., Erdfelder, E., Buchner, A., & Lang, A. G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, 41, 1149–1160. <https://doi.org/10.3758/brm.41.4.1149>
- Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives*, 19(4), 25–42. <https://doi.org/10.1257/089533005775196732>
- Frankfurt, H. (2005 [1986]). *On bullshit*. Princeton, NJ: Princeton University.
- Frankfurt, Harry G. 2006. *On truth*. New York: Alfred A. Knopf.
- Gligorić, V., Feddes, A., & Doosje, B. (2020, October 27). Political bullshit receptivity and its correlates: a cross-cultural validation of the concept. <https://doi.org/10.31234/osf.io/u9pe3>
- Grant, A. M., Franklin, J., & Langford, P. (2002). The Self-Reflection and Insight Scale: A new measure of private self-consciousness. *Social Behavior and Personality: An international journal*, 30(8), 821-836. <https://doi.org/10.2224/sbp.2002.30.8.821>
- Handel, M. I. (1982). Intelligence and deception. *The Journal of Strategic Studies*, 5(1), 122-154. <https://doi.org/10.1080/01402398208437104>

- Higgins, E. T. (2012). Regulatory focus theory. In P. A. M. Van Lange, A. W. Kruglanski, & E. T. Higgins (Eds.), *Handbook of theories of social psychology* (p. 483–504). Sage Publications Ltd. <https://doi.org/10.4135/9781446249215.n24>
- Hudson, C. A., Vrij, A., Akehurst, L., Hope, L., Satchell, L. P. (2020). Veracity is in the eye of the beholder: A lens model examination of consistency and deception, *Applied Cognitive Psychology*, 34(5), 996-1004. <https://doi.org/10.1002/acp.3678>
- Jerrim, J., Parker, P., & Shure, D. (2019). Bullshitters: Who are they and what do we know about their lives? (April 2019). *IZA Discussion Paper No. 12282*. Available at SSRN: <https://ssrn.com/abstract=3390272>
- Kristansen, L. J., & Kaussler, B. (2018). The Bullshit Doctrine: Fabrications, lies, and nonsense in the age of Trump. *Informal Logic* 38(1), 13–52. doi:10.22329/il.v38i1.5067. Krumrei-
- Mancuso, E. J., & Rouse, S. V. (2016). The development and validation of the Comprehensive Intellectual Humility Scale. *Journal of Personality Assessment*, 98, 209-221. <https://doi.org/10.1080/00223891.2015.1068174>
- Litman, L., Robinson, J., & Abberbock, T. (2016). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 49(2), 1 - 10. <https://doi.org/10.3758/s13428-016-0727-z>
- Littrell, S., Fugelsang, J. A., & Risko, E. F. (2020). Overconfidently underthinking: narcissism negatively predicts cognitive reflection, *Thinking & Reasoning*, 26(3), 352-380. <https://doi.org/10.1080/13546783.2019.1633404>
- Littrell, S., Risko, E. F., & Fugelsang, J. A. (2021). The Bullshitting Frequency Scale: Development and psychometric properties. *British Journal of Social*

Psychology, 60(1), 248-270. <https://doi.org/10.1111/bjso.12379>

Lipkus, I. M., Samsa, G., and Rimer, B. K. (2001). General performance on a numeracy scale among highly educated samples. *Medical Decision Making*, 21, 37-44.

<https://doi.org/10.1177/0272989x0102100105>

Malhotra, N., Krosnick, J. A., & Haertel, E. (2007). *The Psychometric Properties of the GSS Wordsum Vocabulary Test (GSS Methodology Report No. 111)*. Chicago: NORC.

McCarthy, I. P., Hannah, D., Pitt, L. F., & McCarthy, J. M. (2020). Confronting indifference toward truth: Dealing with workplace bullshit. *Business Horizons*, 63(3), 253-263.

<https://doi.org/10.1016/j.bushor.2020.01.001>

Mears, D. (2002). The ubiquity, functions, and contexts of bullshitting. *Journal of Mundane Behavior*, (3)2, 233-256. Retrieved from

https://www.researchgate.net/publication/289724155_The_ubiquity_functions_and_contexts_of_bullshitting

Meibauer, J. (2016). Aspects of a theory of bullshit. *Pragmatics & Cognition*, 23(1), 68-91.

<https://psycnet.apa.org/doi/10.1075/pc.23.1.04mei>

Michels, M., Molz, G., & Maas genannt BERPohl, F. (2020). The ability to lie and its relations to the dark triad and general intelligence. *Personality and Individual Differences*, 166,

110195. <https://doi.org/10.1016/j.paid.2020.110195>

Pennycook, G., Cheyne, J. A., Barr, N., Koehler, D. J., & Fugelsang, J. A. (2015). On the reception and detection of pseudo-profound bullshit. *Judgment and Decision Making*,

10(6), 549–563. <http://journal.sjdm.org/15/15923a/jdm15923a.pdf>

- Pennycook, G., McPhetres, J., Bago, B., & Rand, D. G. (2020, April 14). Attitudes about COVID-19 in Canada, the U.K., and the U.S.A.: A novel test of political polarization and motivated reasoning. *PsyArxiv*, <https://doi.org/10.31234/osf.io/zhjkg>
- Pennycook, G., & Rand, D. G. (2020). Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *Journal of Personality*, 88(2), 185-200. <https://doi.org/10.1111/jopy.12476>
- Pennycook, G., Ross, R. M., Koehler, D. J., & Fugelsang, J. A. (2017). Dunning-Kruger effects in reasoning: Theoretical implications of the failure to recognize incompetence. *Psychonomic Bulletin & Review*, 24(6), 1774–1784. <https://doi.org/10.3758/s13423-017-1242-7>
- Peter, J., Churchill, G., & Brown, T. (1993). Caution in the Use of Difference Scores in Consumer Research. *Journal of Consumer Research*, 19(4), 655-662. <https://doi.org/10.1086/209329>
- Primi, C., Morsanyi, K., Chiesi, F., Donati, M. A., & Hamilton, J. (2016). The development and testing of a new version of the Cognitive Reflection Test applying item response theory (IRT). *Journal of Behavioral Decision Making*, 29(5), 453-469. <https://doi.org/10.1002/bdm.1883>
- Pronin, E., Lin, D. Y., & Ross, L. (2002). The Bias Blind Spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin*, 28(3), 369–381. <https://doi.org/10.1177/0146167202286008>
- Reisch, G. A. (2006). “The pragmatics of bullshit, intelligently designed.” In G. L. Hardcastle & G. A. Reisch (Eds.), *Bullshit and philosophy: Guaranteed to get perfect results every time* (pp. 33–47). Chicago: Open Court.

Spicer, A. (2020). Playing the bullshit game: How empty and misleading communication takes over organizations. *Organization Theory*, 1(1), 1-26.

<https://doi.org/10.1177%2F2631787720929704>

Thomson, K. S., & Oppenheimer, D. M. (2016). Investigating an alternate form of the cognitive reflection test. *Judgment and Decision Making*, 11(1), 99–113.

Thorndike, R. L. (1942). Two screening tests of verbal intelligence. *Journal of Applied Psychology*, 26(2), 128–135. <https://doi.org/10.1037/h0060053>

Turpin, M. H., Walker, A. C., Kara-Yakoubian, M., Gabert, N. N., Fugelsang, J. A., & Stolz, J. A., (2019). Bullshit makes the art grow profounder. *Judgment and Decision Making*, 14(6), 658-670. <http://journal.sjdm.org/19/190712/jdm190712.html>

Turpin, M. H., Kara-Yakoubian, M., Walker, A. C., & Fugelsang, J. A. (2020, May 11). Bullshit ability as an honest signal of intelligence. *PsyArxiv*, <https://doi.org/10.31234/osf.io/aru3f>

Vickers, A. J., & Altman, D. G. (2001). Analysing controlled trials with baseline and follow up measurements. *BMJ*, 323, 1123-1124. <https://doi.org/10.1136/bmj.323.7321.1123> Walker, A. C.,

Turpin, M. H., Stolz, J. A., Fugelsang, J. A., & Koehler, D. J. (2019). Finding meaning in the clouds: Illusory pattern perception predicts receptivity to pseudo-profound bullshit. *Judgment and Decision Making*, 14(2), 109-119.

Webber, Jonathan (2013). Liar! *Analysis*, 73(4), 651–659.

<https://doi.org/10.1093/analys/ant081>

Wright, G. R. T., Berry, C. J., & Bird, G. (2012). “You can’t kid a kidder”: association between production and detection of deception in an interactive deception task. *Frontiers in*

Human Neuroscience, 6(87), 1-7. <https://doi.org/10.3389/fnhum.2012.00087>

Wright, G. R. T., Berry, C. J., & Bird, G. (2013). Deceptively simple...The “deception-general” ability and the need to put the liar under the spotlight. *Frontiers in Neuroscience*, 7(152), 1-9. <https://doi.org/10.3389/fnins.2013.00152>

Zvi L., Elaad E. (2018). Correlates of narcissism, self-reported lies, and self-assessed abilities to tell and detect lies, tell truths, and believe others. *Journal of Investigative Psychology and Offender Profiling*, 15(3), 271–286. <https://doi.org/10.1002/jip.1511>

