The psychology of bullshitting: Measurement, correlates, and outcomes of the propensity to mislead others

by

Shane Littrell

A thesis

presented to the University of Waterloo

in fulfillment of the

thesis requirement for the degree of

Doctor of Philosophy

in

Cognitive Psychology

Waterloo, Ontario, Canada, 2021

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner	Philip Fernbach Associate Professor Leeds School of Business University of Colorado – Boulder
Supervisors	Jonathan Fugelsang Professor Department of Psychology University of Waterloo
	Evan Risko Associate Professor Department of Psychology University of Waterloo
Internal Members	Derek Koehler Professor Department of Psychology University of Waterloo
	Ramona Bobocel Professor Department of Psychology University of Waterloo
Internal-external Member	John McLevey Associate Professor Department of Knowledge Integration University of Waterloo

Author's declaration

This thesis consists of material all of which I authored or co-authored (see Statement of Contributions included in the thesis). This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

The present research has been published in the *British Journal of Social Psychology*.

Chapters 1 and 2:

A version of the research presented in Chapters 1 and 2 appears in:

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

Chapter 3:

A version of the research presented in Chapter 3 appears in:

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,* Advance online publication. https://doi.org/10.1111/bjso.12447

Abstract

Recent psychological research has identified important individual differences associated with receptivity to bullshit, which has greatly enhanced our understanding of the processes behind susceptibility to pseudo-profound or otherwise misleading information. However, the bulk of this research attention has focused on cognitive and dispositional factors related to bullshit (the product), while largely overlooking the influences behind bullshitting (the act). Here, I present results from nine studies focusing on: 1) the construction and validation of a new, reliable scale measuring the frequency with which individuals engage in two types of bullshitting (persuasive and evasive) in everyday situations; 2) the associations of both types of bullshitting frequency with other relevant constructs, and; 3) the extent to which those who produce bullshit are also receptive to various types of bullshit. Overall, bullshitting frequency was negatively associated with sincerity, honesty, cognitive ability, open-minded cognition, and self-regard. Additionally, the Bullshitting Frequency Scale was found to reliably measure constructs that are (1) distinct from lying and (2) significantly related to performance on overclaiming and social decision tasks. Moreover, the frequency with which individuals engage in persuasive bullshitting (i.e., bullshitting intended to impress or persuade others) was found to positively predict susceptibility to various types of misleading information and this association is robust to individual differences in cognitive ability and analytic cognitive style. These results represent an important step forward in the study of the spread of misinformation by demonstrating the utility of the Bullshitting Frequency Scale as well as highlighting certain individual differences that may play important roles in the extent to which individuals engage in and are receptive to everyday bullshitting.

Acknowledgements

I would like to express my deepest appreciation and thanks to my supervisors, Jonathan Fugelsang and Evan Risko, for their expert mentorship, patience, and thoughtful guidance throughout my research experience at University of Waterloo. I am a much better researcher, scientist, and person for having the opportunity to work with them and get to know them over the past few years. I could not have asked for better supervisors and consider myself very lucky to have been granted such a wonderful opportunity to learn and grow. Also, many thanks and sincere appreciation go to the mentors I had during my master's program, Chris Cunningham, Amye Warren, and Jill Shelton. If not for your guidance, instruction, and help, I would never have made it this far.

I would also like to thank my lab mates and colleagues for their support and friendship. Were it not for Ethan Meyers, Martin Turpin, and Alex Walker to bounce ideas off of, my research would have been far less interesting and, more importantly, my life would have been filled with far fewer laughs. To Michelle Ashburner, Laura Bianchi, Connor Gaspar, Megan Kelly, Daev McLean, April Pereira, Kaiden Stewart, Lu Xinyi, Torin Young, and Mona Zhu, I could not have made it this far without your help, support, friendship, jokes, sarcasm, and your willingness to pilot test my surveys (ha!).

Dedication

This is dedicated to my parents, Steve and Jackie, and my sister, Michelle. I would be nowhere without your love and support.

Table of Contents

List of Figures

List of Tables

Chapter 1 – Introduction

"One of the most salient features of our culture is that there is so much bullshit. Everyone knows this. Each of us contributes his share. But we tend to take the situation for granted. Most people are rather confident of their ability to recognize bullshit and to avoid being taken in by it. So, the phenomenon has not aroused much deliberate concern, or attracted much sustained inquiry. In consequence, we have no clear understanding of what bullshit is, why there is so much of it, or what functions it serves. And we lack a conscientiously developed appreciation of what it means to us. In other words, we have no theory." – Harry Frankfurt (2005|1986)

Given the increasing prevalence of misleading information and "fake news" on the internet and throughout society at large (Pennycook & Rand, 2019), a growing body of work has emerged that focuses on better understanding the nature of *bullshit* and *bullshitting*. Some has been more descriptive, highlighting the use of bullshitting in politics (Kristansen & Kaussler, 2018; Mears, 2002), business organizations (Martin & Wilson, 2011; Spicer, 2013), academic settings (Cohen, 2012), and everyday life (Frankfurt, 1986). Other research has taken a more empirical approach, examining individual differences associated with *receptivity to bullshit*, such as its relations to analytic thinking and biased pattern perception (Pennycook, Cheyene, Barr, Koehler, & Fugelsang, 2015; Walker, Turpin, Stolz, Fugelsang, & Koehler, 2019). Additionally,

recent work has sought to examine the instrumental functions of *bullshitting* as a strategy for managing impressions and attitude change across a broad range of social interactions (Mears, 2002; Petrocelli, 2018).

1.1 Bullshitting, broadly defined

Philosopher Harry Frankfurt (1986) is perhaps best known for his seminal piece, *On Bullshit*, in which he described a "bullshitter" as a person who deliberately conveys a false/phony impression of himself or his intentions in a way that is "disconnected from a concern with the truth" (p. 12). The bullshitter is not necessarily being intentionally untruthful, according to Frankfurt, but he is certainly "faking things." Frankfurt contrasts this from lying in that the liar knows the truth but is deliberately attempting to get others to believe a falsehood (Frankfurt, 1986; also see Hart, Jones, & Terrizzi, 2019).

However, as some have pointed out, it is arguable whether one can have a "misrepresentational intent" (Meibauer, 2018) – that is, intentionally faking or misleading – while simultaneously being completely unconcerned with the truth. For instance, Stokke and Fallis (2017) instead characterize bullshitting as speech that has a "loose concern" with the truthful advancement of conversational progress, rather than simply unconcerned with the veracity of each statement. Indeed, the bullshitter may not actually know the truth-value of every statement he makes, yet he is often *aware of his unawareness*, and asserts himself with a sense of certainty that the totality of his statements is true regardless (Meibauer, 2018). Given this, rather than being completely "unconcerned with the truth," it might be more accurate to say instead that the bullshitter is *epistemically insouciant*, showing the truth a casual, loose concern or indifference (Cassam, 2018; Stokke & Fallis, 2017). Additionally, as Reisch (2006) has pointed

out, what makes some statements "bullshit" is not necessarily the speaker's casual (dis)regard for the truth, but more in the "uses and purposes" for which they employ bullshit. For example, Kimbrough (2006) argues that bullshitting is useful in situations where directness might be impolite or hurt another's feelings and that it is this facet of bullshitting – i.e., saving others from pain – that is more important to its definition than Frankfurt's notion of a lack of concern for the truth. Ultimately, the veracity of what the bullshitter says does not matter to him nearly as much as his motivations for saying it (Cohen, 2012; Mears, 2002; Reisch, 2006).

Building from Frankfurt's work, Mears (2002) more specifically defined bullshitting as a type of communication aimed at creating or maintaining "misleading, yet possible, though frequently improbable, accounts or impressions of self or reality" (p. 236). That is, bullshitting employs rhetoric such as exaggerations, embellishment, and joking in an attempt to manage self-image by presenting oneself in an exaggerated positive light, such as being more competent, intelligent, skilled, or moral than perhaps one actually is or believes himself to be. This definition implicitly rejects the "lack of concern for the truth" element of Frankfurt's definition, echoing Reisch's (2006) notion that bullshitters pragmatically utilize hyperbolic, yet often inconsequential, claims to promote certain concerns, goals, or agendas. Under this view, bullshitting is both instrumental and performative, in that it is a strategy employed to help boost one's self-concept and better navigate and/or gain advantage in a range of social contexts (Mears, 2002).

Others have further veered away or expanded on the Frankfurtian definition of bullshitting, as it captures "just one flower in the lush garden of bullshit" (Cohen, 2012). For instance, Cohen (2012) emphasized that the aim of some bullshitters is to impress using

discourse constructed with "*unclarifiable unclarity*"; that is, relying on vacuous, confusing buzzwords which obscure that the statements, while superficially impressive, contain no discernible meaning (e.g., jargon-heavy writing found in some academic publications). For Cohen, this type of bullshitting is distinct from the purely Frankfurtian type in that the Cohen-bullshitter is unconcerned with the lucidity of what he says, rather than simply unconcerned with its truth-value (Cohen, 2012). The key to bullshitting, then, is to impress or otherwise mislead using an impressively worded yet impenetrably obscure message. In this way, the foundational element of bullshit/bullshitting is a lack of concern for conveying substantive meaning, rather than a lack of concern for the truth.

Additionally, Carson (2016) pointed out that people also sometimes engage in *evasive bullshitting*, a type of *digressive circumlocution* employed in an attempt to simultaneously avoid lying while also avoiding directly answering questions one does not want to answer, due to having insufficient information and/or because giving a direct answer may cause harm to oneself or others (Carson, 2016). In this way, the evasive bullshitter can pragmatically and strategically avoid being untruthful, per se, with less risk of reputational harm, by being "slippery" when navigating socially precarious interactions (Carson, 2016; Mears, 2002; Reisch, 2006). For instance, a politician may be motivated to engage in evasive bullshitting when questioned by a member of the press (see Cillizza, 2019) if, for instance, a direct answer could potentially cost votes (harm to self) or jeopardize national security (harm to others). In some respects, this is similar to the concept of prosocial lying (i.e., lying to benefit or prevent harm to others), except that prosocial lying requires responding with an untruth (lie), whereas *evasive bullshitting* is an attempt to avoid lying, often by substituting a non-relevant truth for a direct response (Carson, 2016; also see, Lupoli, Jampol, & Oveis, 2017). In this way, the evasive bullshitter can save face

or spare feelings by talking around the question to, in essence, "answer without answering." This again underscores the key distinction; a liar's goal is to craft false beliefs in others whereas a bullshitter's goal is to foster or maintain positive impressions, or at least avoid negative ones (Hart, et al., 2019; Mears, 2002).

1.2 Bullshitting as a topic of psychological study

Based on Frankfurt's (1986) work, Pennycook and colleagues (Pennycook et al., 2015) introduced the Bullshit Receptivity Scale (BSR), comprising vacuous, yet grammatically correct, buzzword-heavy statements randomly generated by a computer algorithm[1]. Participants rate each statement according to its perceived profoundness. Higher scores indicate that a person is more receptive to "pseudo-profound bullshit" and have been found to be associated with decreased engagement in reflective thinking (Pennycook et al., 2015), illusory pattern perception (Walker et al., 2019), greater susceptibility to fake news (Pennycook & Rand, 2019), and stronger beliefs in alternative medicine and paranormal phenomena (Čavojová, Secară, Jurkovič, & Šrol, 2019; see also Erlandsson, Nilsson, Tinghög, & Västfjäll, 2018).

While people encounter various forms of bullshit in their daily lives, they also produce their own bullshit. Though *bullshitting* functions across a broad range of social interactions in everyday life, as Mears (2002) noted, this ubiquitous social phenomenon has thus far received little research attention. Recent work by Petrocelli (2018) represents a shift toward empirically examining this common mode of discourse by focusing on *bullshitting* (the act) rather than

[1] The BSR items are arguably examples of combined Cohen-Frankfurt bullshit in that they use vacuous statements crafted with both "unclarifiable unclarity" (Cohen), in the form of pseudo-profound buzzwords, and a loose concern for the truth (Frankfurt), in that they were randomly assembled via computer algorithm.

bullshit (the product). In his study, participants were given opportunities to engage in bullshitting by writing summaries on a given topic that they were told would be evaluated by either an expert or a non-expert. Results supported some Frankfurtian notions regarding bullshitting, suggesting that participants were more likely to engage in bullshitting on those tasks where they felt more *obligated to provide an opinion* and in situations where they felt bullshitting would be *easier to get away with* (Petrocelli, 2018). Though not an exhaustive list of bullshitting antecedents, these results do represent a solid first step in the empirical study of engagement in bullshitting. However, to facilitate advancement in this burgeoning area of research, convenient, standardized measurements and methods will need to be developed.

Chapter 2

Development and psychometric properties of The Bullshitting Frequency Scale

2.1 Present investigation

The following four studies focus on the construction and validation of the Bullshitting Frequency Scale (BSF), a new tool designed to measure the frequency with which individuals engage in "everyday bullshitting," broadly defined. In Study 1, I report the initial development and factor analysis of the scale using items based on definitions taken from philosophical and linguistic literature on bullshitting. This is followed by a series of studies further developing the scale and deepening our understanding of the frequency with which people engage in bullshitting.

2.2 Study 1 – Scale creation

Study 1 served two general goals: 1) initial creation of the scale, and; 2) examination of associations of the new scale with theoretically-related constructs. With respect to (1), I first generated a list of items based on past literature on bullshitting, then administered these to a large sample and used both exploratory and confirmatory analyses to realize the final scale. These procedures and results are presented in Study 1a. With respect to (2), using the same sample, I then examined bivariate and partial associations of the scale with various individual difference measures of related constructs, the results of which I will present in Study 1b.

2.2.1 Method

2.2.1.1 Participants

In order to achieve a sufficient sample size for all analyses, three hundred ninety-one participants from the United States and Canada were recruited via Amazon's Mechanical Turk using the CloudResearch crowdsourcing platform (Litman, Robinson, & Abberbock, 2016). Though there is no universally agreed upon standard for how many participants to recruit for exploratory factor analyses, several helpful criteria have been proposed in the literature. For instance, Nunally (1978) suggested that a sample ten times the number of variables or up to 300 participants is adequate, while Bryant and Yarnold (1995) proposed a subject-to-variable ratio of 5 (which, in the present case, would be 90). The final sample size exceeded both of these recommendations.

Data were collected across two samples (June 2019 and January 2020) and combined into one data set. All raw data files can be found at https://osf.io/dh6vj/. Sixteen participants were removed from the data set for failing attention checks. An additional 14 were removed for receiving a score of less than 0.5 from Google's reCAPTCHA v3 "bot detection" feature (suggesting the responses were likely submitted by a computer algorithm, i.e., "bots"), or for providing notably unusual comments to open-ended numeracy/math problems (e.g., responding with "yes good and nice survey" or copying/pasting the question as the answer), based on recommendations from Chmielweski and Kucker (2019). This left data for 361 participants in the final analysis (222 male, 137 female, 2 intersex or prefer not to answer, M_{age} = 36.40, SD_{age} = 11.26, Bachelor's degree or higher = 50.2%), which provided .90 power to detect an effect of r = .20 at an alpha = .01 (g*power; Faul, Erdfelder, Buchner, & Lang, 2009). Participation was

restricted to those who had at least a 95% MTurk HIT (Human Intelligence Task) approval rating and had completed a minimum of 100 surveys. Participants were paid $3.00 USD for their time.

2.2.1.2 Procedure

After indicating consent and answering demographic questions (i.e., age, biological sex, and level of education), participants were presented with 18 items in randomized order describing various scenarios (based on definitions from previously discussed literature) in which a person might be tempted to engage in bullshitting, As the scale items were designed to capture "everyday bullshitting," broadly construed, there were no *a priori* expectations regarding factor structure. Participants were asked to rate on a 5-point frequency scale from "Never" to "A lot / All the time" how often, in general, they engage in bullshitting as described in each item. Higher scores are meant to indicate that a person reports engaging in bullshitting more frequently. The terms "bullshit" and "bullshitting" were not included in the instructions or scale items. Full instructions given to participants can be found in the supplementary materials.

2.2.2 Results

Based on recommendations from Kim (2013), one item ("Regardless of whether I actually know what I'm talking about") was removed for having skewness with a high absolute z-score value, z(skew) = 6.77. Data for the remaining 17 items were analysed using exploratory principal axis factoring with oblique rotation (direct oblimin), as it was believed that any possible distinct factors of bullshitting that might emerge would be both conceptually and statistically related (Table 1). Sampling adequacy was confirmed via the Kaiser-Meyer-Olkin procedure, yielding a KMO score of .95, which is well above Kaiser's (1974) minimum acceptable level of .50 and exceeds the "marvellous" threshold of .90 proposed by Hutcheson

and Sofroniou (1999). Two factors emerged with eigenvalues above Kaiser's (1974) suggested cut-off criterion of 1.0. Factor 1 had an eigenvalue of 9.13 and accounted for 53.69% of the variance. Factor 2 had an eigenvalue of 1.11 and accounted for 6.54% of the variance. Further analysis of both the rotated component plot and the scree plot justified a two-factor solution as best representing the data (Zwick & Velicer, 1982). Results for this initial factor analysis are show in Table 1.

Table 1

Pattern matrix factor loadings for all 18 items after rotation for each scale item

		1	2	*M*	*SD*
1	When I want to impress the people I'm talking to.	.871		2.50	1.09
2	When I want others to see me as more intelligent or knowledgeable.	.850		2.44	1.02
3	When I know it will be easy to get away with it	.801		2.27	1.04
4	When I know it will help me achieve a goal.	.783		2.53	1.09
5	When I know it will get me what I need or want.	.756		2.48	1.07
6	When I'm trying to fit in better or be more accepted by the person or people I'm interacting with.	.731		2.36	1.07
7	When I want the thing(s) I'm talking about to sound more interesting or exciting.	.723		2.65	1.00
8	When I want to contribute to a conversation or discussion even though I'm not well-informed on the topic.	.703		2.29	1.05
9	By pretending to know more about a topic than I actually do.	.648		2.42	0.93
10	When I feel obligated to share my opinion.	.489		2.44	1.06
11	When I'm trying to avoid looking stupid.	.481	.301	2.60	1.09
12	When I'm "put on the spot" and asked about something I don't know much about.	.458		2.50	1.04
13	When I don't want to tell someone what I really think.		.882	2.65	1.06
14	When someone asks me something that I want to avoid giving a direct answer to.		.661	2.59	1.05
15	When being fully honest would be harmful or embarrassing to me or someone else.		.570	2.91	1.08
16	When I need to fake/bluff my way out of a conversation or situation that I don't want to be in.		.516	2.62	1.01
17	When I want to deflect criticism or questions that might make me look bad.	.360	.461	2.46	1.08
18	Regardless of whether I know what I'm talking about.	*(removed)*			

Extraction Method: Principal Axis Factoring. *Rotation Method*: Oblimin with Kaiser Normalization. Rotation converged in 6 iterations.

To reduce scale size and ensure high factor reliability, my next step was to eliminate redundant items and those with factor loadings < .500. A second principal axis factor analysis was conducted which yielded the same pattern of factor loadings, therefore this 12-item iteration was retained as the final version of the scale (Table 2). Macdonald's scale reliabilities were strong for the full, 12-item scale (ω = .93).

Table 2

Pattern matrix factor loadings for 12-item scale after rotation for each scale item

		1	2	*M*	*SD*
1	When I want others to see me as more intelligent or knowledgeable.	.869		2.44	1.02
2	When I want to impress the people I'm talking to.	.848		2.50	1.09
3	When I know it will be easy to get away with it.	.745		2.27	1.04
4	When I want to contribute to a conversation or discussion even though I'm not well-informed on the topic.	.743		2.29	1.05
5	When I'm trying to fit in better or be more accepted by the person or people I'm interacting with.	.723		2.36	1.07
6	When I know it will help me achieve a goal.	.700		2.53	1.09
7	By pretending to know more about a topic than I actually do.	.691		2.42	0.93
8	When I want the thing(s) I'm talking about to sound more interesting or exciting.	.691		2.65	1.00
9	When I don't want to tell someone what I really think.		.827	2.65	1.06
10	When someone asks me something that I want to avoid giving a direct answer to.		.722	2.59	1.05
11	When being fully honest would be harmful or embarrassing to me or someone else.		.657	2.91	1.08
12	When I need to fake/bluff my way out of a conversation or situation.		.534	2.62	1.01

Extraction Method: Principal Axis Factoring. Rotation Method: Oblimin with Kaiser Normalization. Rotation converged in 4 iterations. Factor correlation, $r = .76$

2.2.2.1 Confirmatory analysis

I next conducted a confirmatory factor analysis (CFA) using JASP (v0.11.1.0) to confirm whether a two-factor structure was a better fit for the data. Results confirmed that the two-factor model ($\chi^2(53) = 128.59$, $p < .01$; CFI = .97; TLI = .96; RMSEA = .06) was a better fit to the data compared to a one-factor model ($\chi^2(54) = 243.88$, $p < .01$; CFI = .92; TLI = .90; RMSEA = .10). Figure 1 shows factor loading plots for both the one-factor and two-factor model. Fit indices for both models are listed in Table 3.

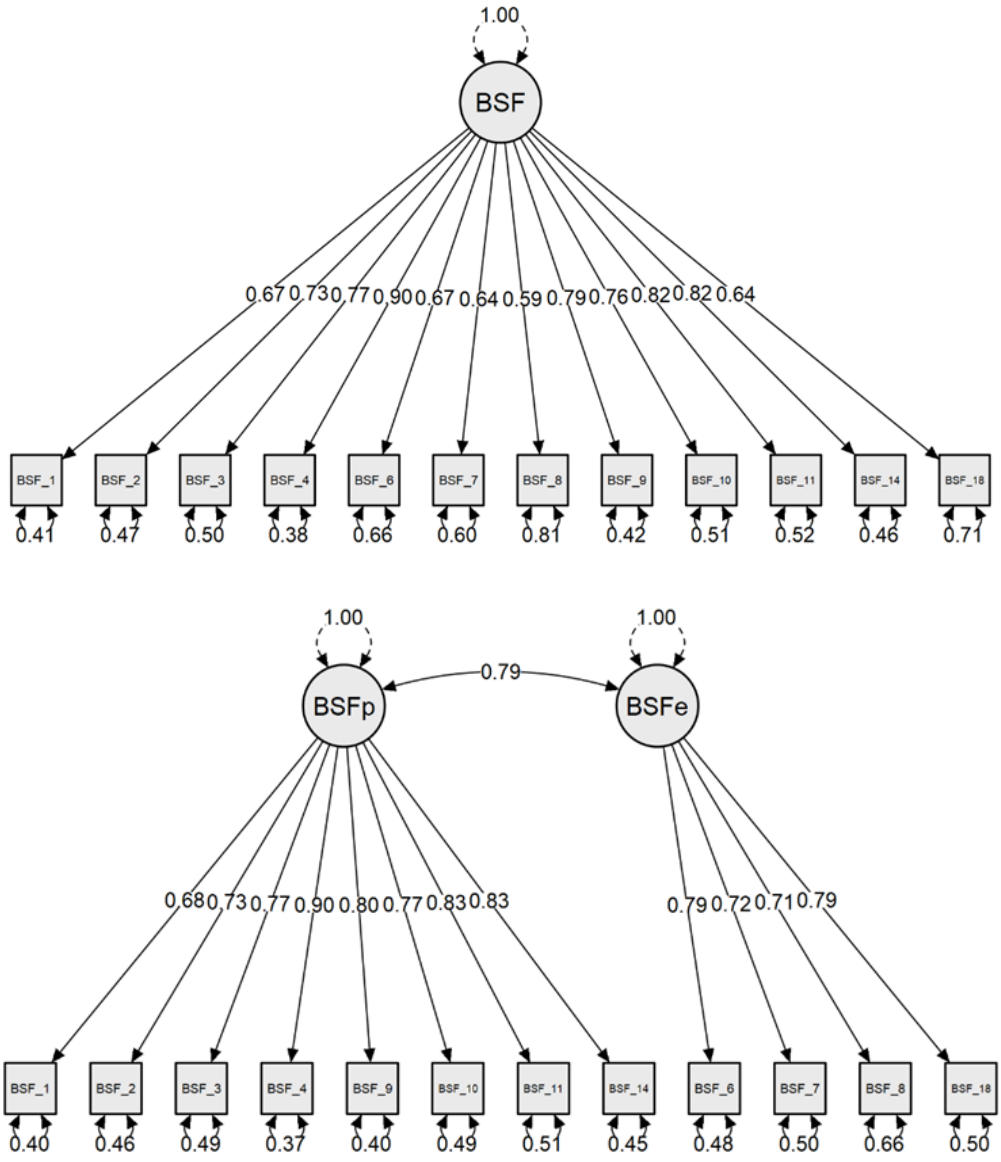

Figure 1. Factor loading plots for one-factor (top) and two-factor (bottom) models of the 12-item Bullshitting Frequency Scale. BSFp = Persuasive bullshitting; BSFe = Evasive bullshitting

Table 3
Fit indices for one-factor and two-factor BSF models

Index	1 Factor	2 Factor
Comparative Fit Index (CFI)	0.92	0.97
Tucker-Lewis Index (TLI)	0.90	0.96
Bentler-Bonett Non-normed Fit Index (NNFI)	0.90	0.96
Bentler-Bonett Normed Fit Index (NFI)	0.90	0.95
Parsimony Normed Fit Index (PNFI)	0.73	0.76
Bollen's Relative Fit Index (RFI)	0.87	0.93
Bollen's Incremental Fit Index (IFI)	0.92	0.97
Relative Noncentrality Index (RNI)	0.92	0.97
Additional metrics	1 Factor	2 Factor
Root mean square error of approximation (RMSEA)	0.10	0.06
Standardized root mean square residual (SRMR)	0.05	0.03
Hoelter's critical N (α = .05)	104.85	194.79
Hoelter's critical N (α = .01)	117.68	218.95
Goodness of fit index (GFI)	0.89	0.95
McDonald fit index (MFI)	0.76	0.90
Expected cross validation index (ECVI)	0.83	0.51

2.2.2.2 Factor labelling

The items clustering around Factor 1 reflect engagement in a type of bullshitting that: 1) is motivated by a desire to impress and be accepted by others; 2) often involves misrepresenting oneself as more intelligent or knowledgeable about a topic than he/she actually is; 3) can include language meant to be perceived as superficially interesting or exciting, and; 4) is enacted when perceived to be easy to get away with. These elements appear to align most closely with various aspects of bullshitting as defined by Frankfurt (1986), Mears (2002), Cohen (2012), Reisch (2006), and Stokke and Fallis (2017), and I have labelled this factor *persuasive bullshitting.* Cronbach's alpha scale reliability for the *persuasive* factor was strong (α = .92).

The items clustering around Factor 2 suggest that it represents bullshitting initiated when a person does not want to reveal what he/she thinks about a particular topic, believes that answering a question(s) in a frank manner would be harmful or embarrassing, and/or wants to avoid an inquiry altogether. As this description appears to most closely align with Carson's (2016), Kimbrough's (2006), and elements of Mears (2002) views of bullshitting, I have labelled this factor *evasive bullshitting.* Cronbach's alpha scale reliability for this factor was also strong ($\alpha = .81$).

2.2.3 Discussion

The results presented here suggest that the frequency with which people engage in everyday bullshitting can be captured using a self-report measure (i.e., the Bullshitting Frequency Scale) that conceptualizes bullshitting in terms of two main factors. The first factor, deemed *persuasive bullshitting,* involves positively-biased misrepresentations of one's own knowledge, attitudes or skills (Frankfurt, 1986; Mears, 2002) and uses rhetorical tactics including boasting or puffery meant to make oneself or what one is saying seem more interesting, impressive, or otherwise persuasive (Cohen 2012; Frankfurt, 1986; Mears, 2002). This arguably captures the notion of "bullshitting" as it is commonly understood by the general public.

The second factor, *evasive bullshitting*, reflects a strategic evasiveness or bluffing motivated by a desire to avoid giving direct answers to, or otherwise participating in, some inquiry where more direct responses might result in undesirable social costs to self or others (Carson, 2016; Mears, 2002; Stokke & Fallis, 2017). As noted, this can be done for selfish or noble/altruistic reasons (e.g., navigating polite conversation) but, just as prosocial lying (no

matter how altruistically intentioned) is still lying, by definition, evasive bullshitting (even if altruistically intentioned) is still bullshitting, by definition (see Cheung, Siu, & Chen, 2015).

2.3 Study 2 – Associations with related constructs

I next examined correlations between the *Bullshitting Frequency Scale* (BSF) and individual differences measures thought to be conceptually related to bullshitting. As bullshitting is believed to be partially motivated by a loose (or less) concern for the truth of what one is saying (Frankfurt, 1986; Stokke & Fallis, 2017), and misleading representations of "what one is up to" (Frankfurt, 1986; Mears, 2002), it was thought that the BSF would show moderate associations with measures of trait honesty and sincerity. Also, as bullshitting (at least on its face) involves misleading self-descriptions and attempts to give distorted impressions (Carson, 2016; Mears, 2002), it was thought that BSF scores would be associated with measures of social desirability.

Additionally, while it is important for the scale to be correlated with measures of related constructs, it also needs to predict actual behaviour, for instance on a task that previous literature has suggested is arguably an instance of bullshitting (Jerrim, Parker, & Shure, 2019; Pennycook & Rand, 2019). Therefore, participants also completed the Overclaiming Questionnaire (Paulhus et al., 2003). Finally, as past research has separately found that honesty is negatively associated with cognitive ability (Kajonius, 2014; Ruffle & Tobol, 2017), this would suggest that *bullshitting frequency* might be positively associated with cognitive ability. However, research on lying has failed to find significant associations with cognitive ability, which appears inconsistent with the results from the honesty literature (Wright, Berry, Bird, 2012; Wright,

Berry, Bird, 2013). Therefore, to examine bullshitting frequency's potential associations with cognitive ability, participants also completed measures of numeracy and verbal intelligence.

2.3.1 Method

Study 2 utilized the same sample as Study 1, therefore participants and procedures are identical to those reported earlier for Study 1.

2.3.1.1 Materials

Participants completed the following measures in randomized order (full descriptions of each measure can be found in the supplementary materials):

Bullshitting Frequency Scale

Scores for the 8-item *persuasive* and 4-item *evasive bullshitting* subscales were calculated by computing the mean score for each subscale. An "overall bullshitting" score was then calculated by adding the two subscale means and dividing by two to compensate for the asymmetry in the number of items for each.

Honesty

Honesty was assessed using the Integrity/Honesty/Authenticity scale from the IPIP version of the Values in Action scale (Goldberg et al., 2006; Peterson & Seligman, 2004). Participants rated statements assessing their propensity for behaving with honesty and integrity such as, "I lie to get myself out of trouble" (reverse scored), according to a scale from 1 (strongly disagree) to 5 (strongly agree). Higher scores indicate a higher tendency toward comporting oneself with honesty and integrity.

Sincerity

The extent to which a person represents himself/herself in a sincere way was measured using the IPIP's version of the HEXACO sincerity scale (Goldberg et al., 2006). Participants rated their agreement with statements such as, "(I) use flattery to get ahead," on a scale from 1 (strongly disagree) to 5 (strongly agree). Items whose wordings were redundant with BSF items were removed (e.g., "I play a role in order to impress people "), leaving a 7-item scale.

Overclaiming

The extent to which a person claims to know more about a topic than he/she actually does was measured using an adapted, 30-item version of the overclaiming questionnaire (OCQ; Paulhus et al. 2003). Participants rated their familiarity with items from two lists (one related to historical names/events and the other covering physical sciences) using a scale from 0 (Never heard of it) to 6 (Very familiar). Both lists contained 15 items, 3 of which were fake. Responses were recoded so that indications of any level of familiarity were given a "1" and all those rated as "Never heard of it" were scored as "0." Based on recommendations from Paulhus et al. (2003), and to ensure that the direction of associations were congruent with bullshitting frequency, an overclaiming accuracy score was computed by subtracting the proportion of hits (i.e., familiarity with the genuine/real items) from the proportion of false alarms (i.e., indicating familiarity with fake items). Positive scores indicate a tendency to claim knowledge of items that do not exist. I also calculated total false alarms as a second measure of overclaiming.

Cognitive ability

Numeracy was measured using a 10-item version of the General Risk and Numeracy Scale, which assesses a person's ability to perform and understand basic mathematical operations

(Lipkus, Samsa, & Rimer, 2001). When originally validated across three studies, the scale was reported to have an acceptable average alpha of .73 (Lipkus et al., 2001). Verbal intelligence was measured using a 10-item version of the "Wordsum" vocabulary test (Thorndike, 1942; Malhotra, Krosnick, & Haertel, 2007). The Wordsum has demonstrated an acceptable Cronbach's alpha reliability of $\alpha = .71$ in past research (Littrell, Fugelsang, & Risko, 2020).

Social desirability

As bullshitting involves misleading self-descriptions and attempts to give distorted impressions (Carson, 2016; Cohen, 2012; Frankfurt, 1986; Mears, 2002), it was thought that BSF scores would be associated with measures of social desirability, specifically impression management as measured by the Balanced Inventory of Desirable Responding (BIDR; Paulhus, 1991). The BIDR contains two subscales. The self-deceptive enhancement subscale of the BIDR is purported to measure "honest but overly positive" assessments of oneself (Hart, Ritchie, & Hepper, 2015). Likewise, the impression management subscale is thought to measure one's bias toward pleasing others through honest self-representation while avoiding negative impressions (Hart et al., 2015). Participants rated themselves on a 7-point Likert scale ("strongly disagree" to "strongly agree") for all items. Scores for each subscale were calculated by summing item ratings for each scale.

2.3.2 Results

Table 4 lists descriptive statistics as well as bivariate and partial correlation values for all study variables. It should be noted that recent research has raised validity issues with the BIDR which can obscure clear interpretation of associations with other variables (e.g., Müller & Moshagen, 2019). Indeed, a growing body of research has shown that the BIDR subscales are

positively, rather than negatively, related to current measures of trait honesty (de Vries, Zettler, & Hilbig, 2013; de Vries et al., 2018; Müller & Moshagen, 2019; Uziel, 2010) and may be contaminated to the degree that it confounds honesty, Big Five traits, and biased self-presentation (Bensch, Paulhus, Stankov, & Ziegler, 2019; Müller & Moshagen, 2019). This has led some researchers to advise against using the BIDR as a measure of a self-favouring response bias (Müller & Moshagen, 2018). Given the issues raised in the literature, as well as the present results affirming positive rather than negative relations with honesty-related variables, I encourage caution when interpreting any associations with BIDR variables reported here.

Table 4
Descriptive and correlational data for all study variables with BSF, BSFp, and BSFe

		M	*SD*	α	*Bivariate*			*Partial*	
					BSF	BSFp	BSFe	BSFp[a]	BSFe[b]
1	BSF (overall)	2.55	0.76	*.92*	-			-	-
2	Persuasive bullshitting (BSFp)	2.42	0.83	*.92*	.92**	-		-	-
3	Evasive bullshitting (BSFe)	2.68	0.84	*.82*	.92**	.68**	-	-	-
4	Overclaiming	-0.39	0.30	-	.19**	.23**	.12*	.20**	-.05
5	Overclaiming - false alarms	2.20	1.92	*.77*	.26**	.30**	.18 *	.24**	-.03
6	Honesty	4.15	0.65	*.82*	-.48**	-.49**	-.38**	-.34**	-.08
7	Sincerity	3.77	0.95	*.89*	-.62**	-.64**	-.50**	-.47**	-.11*
8	Self-deceptive enhancement	84.32	16.03	*.81*	-.26**	-.22**	-.24**	-.08	-.13*
9	Impression management	78.20	19.59	*.85*	- .38**	-.35**	-.35**	-.17**	-.16**
10	Cognitive ability	7.69	1.68	*.79*	-.26**	-.32**	-.16**	-.30**	.09

Note: N = 361. BSF = Bullshitting Frequency Scale; BSFp[a] = Persuasive bullshitting, controlling for evasive; BSFe[b] = Evasive bullshitting, controlling for persuasive. $^{**}p < .01$; $^{*}p < .05$

2.3.2.1 Bivariate and partial correlations

At the bivariate level, overall bullshitting frequency scores (BSF) were significantly and positively related to overclaiming, $r(359) = .19$, $p < .001$, and total false alarms, $r(359) = .26$, $p <$

.001, and significantly and negatively related to honesty, $r(359) = -.48$, $p < .001$, sincerity, $r(359) = -.62$, $p < .001$, and cognitive ability, $r(359) = -.26$, $p < .001$. To explore the BSF's two-factor structure in more detail, I also examined the partial correlations of all variables with each of the BSF subscales (i.e., persuasive and evasive) controlling for the other subscale. *Persuasive bullshitting* scores (controlling for evasive) were significantly and positively related to overclaiming, $r(358) = .20$, $p < .001$, and total false alarms, $r(358) = .24$, $p < .001$, and significantly and negatively related to honesty, $r(358) = -.34$, $p < .001$, sincerity, $r(358) = -.47$, $p < .001$, and cognitive ability, $r(358) = -.30$, $p < .001$. *Evasive bullshitting* scores (controlling for persuasive) were significantly related only to sincerity, $r(358) = -.11$, $p = .04$.

2.3.3 Discussion

Consistent with my predictions, overall bullshitting frequency was negatively associated with trait honesty, sincerity, social desirability, and cognitive ability at the bivariate level. Bullshitting frequency was also positively associated with performance on an overclaiming task. That is, not only were BSF scores significantly related to conceptually-related self-report measures, they were also significantly related to claiming knowledge of things that do not exist, which may reflect the tendency to "bullshit oneself" (Pennycook & Rand, 2019). Overall, these results are consistent with the idea that individuals who more frequently engage in bullshitting are less honest, less sincere, and demonstrate lower cognitive ability than those who bullshit less frequently. Additionally, more prolific bullshitters were also more likely to overclaim when asked to demonstrate their general knowledge, possibly bullshitting themselves as well as others (Jerrim, Parker, & Shure, 2019; Pennycook & Rand, 2019).

2.4 Study 3 – Cognitive correlates of bullshitting

In Study 3, I expand my investigation to examine psychological factors thought to underlie motivations to engage in bullshitting. Specifically, I set out to test the claims that engagement in bullshitting involves processes that are "less deliberative and analytical" (Frankfurt, 1986, p.16) and is motivated by concerns related to self-image (Mears, 2002). For this, I will compare bullshitting frequency to measures of analytic thinking styles (i.e., Need for Cognition, Need for Cognitive Closure, open-minded thinking) and self-regard (i.e., core self-evaluations). It was expected that BSF scores would be positively related to Need for Cognitive Closure, and negatively related to Need for Cognition, open-minded thinking, and self-regard.

2.4.1 Method

2.4.1.1 Participants

Two hundred adult participants from the United States and Canada were recruited from Amazon's Mechanical Turk participant pool with the goal of achieving approximately .80 power to detect an effect of r = .20 at α = .05 (g*power; Faul, Erdfelder, Buchner, & Lang, 2009). Data from eight participants were removed for failing attention checks. This left data for 192 participants to consider in the final analysis (108 male, 93 female, 1 intersex, M_{age} = 35.91, SD_{age} = 10.41, Bachelor's degree or higher = 54.7%).

2.4.1.2 Procedure

Recruitment and participation procedures were identical to Study 2 with the exception that individuals who had participated in Study 2 were restricted from participating in Study 3. In

addition to completing the Bullshitting Frequency Scale (BSF), participants were also presented with the following measures in a randomized order.

2.4.1.3 Materials

Open-minded thinking

Participants completed a 10-item measure of actively open-minded thinking (AOT) designed to assess the extent to which a person believes that having a cognitively flexible, open-minded perspective should be a general social norm (Baron, 2019), were rated on a 5-point Likert scale from "strongly disagree" to "strongly agree." High scores indicate that a person highly values open-minded thinking as a general concept.

As the AOT measures one's attitudes toward open-minded thinking in general, I also wanted to assess a person's individual open-mindedness. Therefore, participants completed the 6-item Open-Minded Cognition Scale (OMC; Price et al., 2015). Items were rated on a 7-point Likert scale from "strongly disagree" to "strongly agree." Higher scores indicate that a person more strongly believes that they, personally, are an open-minded, reflective thinker.

Analytic thinking

The degree to which a person enjoys engaging in intellectually effortful activities was measured using the 20-item Need for Cognition (NFC) subscale of the Rational-Experiential Inventory (REI; Pacini & Epstein, 1999). Preference for intuitive thinking (i.e., "trusting one's gut") was assessed with the Faith in Intuition (FI) subscale. Higher NFC scores are thought to indicate a greater preference for analytic thinking while higher FI scores are thought to indicate greater preference for relying on intuitions.

The degree to which ambiguity motivates a person to require/accept an answer (regardless of the correctness of the answer) so that information processing can be discontinued, thus disrupting analytic processes, was measured using the 15-item short version of the revised Need for Cognitive Closure Scale (NFCC; Roets & Van Hiel, 2011; Webster & Kruglanski, 1994). Higher scores indicate a greater degree of discomfort with decisional ambiguity.

Self-regard

The 12-item Core Self-Evaluations Scale (CSES; Judge, Erez, Bono, & Thorensen, 2003) was used to measure an individual's overall feelings of self-esteem, self-efficacy, emotional stability, and locus of control. Higher scores indicate overall positive feelings of self-regard (i.e., self-worth).

2.4.2 Results

Descriptive statistics and Pearson's *r*-values for all study variables are listed in Table 5. Given the lack of relevant prior research, I made no predictions related to intuitive thinking, but report the results here for a more fulsome model of potential bullshitting correlates.

Table 5

Descriptive and correlational data for all study variables with BSF and its subscales

		M	*SD*	α	*Bivariate*			*Partial*	
					BSF	BSFp	BSFe	BSFp[a]	BSFe[b]
1	BSF (overall)	2.70	0.79	*.93*	-				
2	Persuasive bullshitting (BSFp)	2.61	0.91	*.94*	.93**	-			
3	Evasive bullshitting (BSFe)	2.80	0.81	*.78*	. 91**	.69**	-		
4	Actively open-minded thinking	3.95	0.69	*.82*	-.27**	-.32**	-.17**	-.28**	.07
5	Open-minded cognition scale	5.15	1.06	*.82*	-.39**	-.40**	-.32**	-.26**	-.07
7	Need for Cognitive Closure	4.68	1.11	*.92*	.22**	.21**	.19**	.11	.07
6	Need for Cognition	3.73	0.81	*.94*	-.27**	-.25**	-.24**	-.12	-.10
8	Faith in Intuition	3.15	0.85	*.95*	-.13	-.07	-.17*	.06	-.16*
9	Self-regard	3.41	0.76	*.90*	-.36**	-.32**	-.35**	-.11	-.20**

Note: N = 192. BSF = Bullshitting Frequency Scale; BSFp[a] = Persuasive bullshitting, controlling for evasive; BSFe[b] = Evasive bullshitting, controlling for persuasive. **$p < .01$; *$p < .05$

2.4.2.1 Bivariate and partial correlations

At the bivariate level, overall bullshitting frequency scores were significantly and negatively related to actively open-minded thinking (AOT), $r(190) = -.29$, $p < .001$, open-minded cognition (OMC) , $r(190) = -.40$, $p < .001$, need for cognition, $r(190) = -.26$, p < .001, and self-regard, $r(190) = -.35$, $p < .001$, and significantly and positively related to need for cognitive closure, $r(190) = .22$, $p < .001$, but were not related to Faith in Intuition, $r(190) = -.01$, $p = .13$. The pattern of associations for the BSF subscales (i.e., persuasive and evasive) closely followed the same pattern.

I next created adjusted scores for persuasive and evasive bullshitting to account for any overlapping variance by calculating partial correlations for each bullshitting type (BSFp and BSFe) controlling for the other. Doing so allows us to better understand these associations for individuals who more often engage in one type of bullshitting over the other. Results revealed

that adjusted scores for persuasive bullshitting were significantly and negatively related only to actively open-minded thinking (AOT), $r(189) = -.28$, $p < .001$, and open-minded cognition (OMC) , $r(189) = -.26$, $p < .001$. Additionally, adjusted scores for evasive bullshitting were significantly and negatively related only to Faith in Intuition, $r(189) = -.28$, $p < .001$, and self-regard, $r(189) = -.28$, $p < .001$.

2.4.3 Discussion

In Study 3, I found that individuals who reported bullshitting more frequently also reported being less open-minded, experiencing less enjoyment from engaging in cognitively effortful pursuits (i.e., lower need for cognition), less tolerance for informational and decisional ambiguity (i.e., higher need for cognitive closure), and had lower feelings of self-worth. This lends support to Frankfurt's (1986) and Mears' (2002) assertions that processes related to analytic thinking and self-evaluation underlie one's motivations to engage in bullshitting. Therefore, it could be the case that individuals who are less open-minded (i.e., more cognitively rigid) and have lower feelings of self-worth may engage in bullshitting more frequently than more open-minded people with positive self-worth.

Overall, these results are consistent with the idea that bullshitting often involves strategies meant to impress, persuade, and/or foster positive impressions in others by manipulating aspects of one's self-image (Mears, 2002). A person with low self-regard has negative views of his own worth and competence (Judge et al., 2003) and may feel strongly motivated to exaggerate and embellish personal qualities in a positive way to feel more valued in social interactions (Mears, 2002). Additionally, individuals who are more cognitively rigid and less tolerant of divergent perspectives (i.e., are dispositionally less open-minded) might use

bullshitting as a persuasion strategy to leverage greater acceptance of their beliefs, opinions, and ideas, thus minimizing disagreement and attenuating the need for them to change their own minds, which in turn may also provide some ego-protective benefits.

2.5 Study 4 – Bullshitting versus Lying

To further validate the BSF, in Study 4, I sought to discriminate the measurement of "everyday bullshitting" from that of "everyday lying" while also confirming its factor structure. Therefore, I compared the Bullshitting Frequency Scale to the Lying in Everyday Situations scale (LiES; Hart et al., 2019) both at the factor/item level as well as at the bivariate level in terms of associations with other variables of interest. The LiES scale (Hart et al., 2019) was designed as a reliable and valid measure of the *propensity to lie across various contexts in everyday life*. Scores have been found to positively correlate with other popular measures of lying, lie acceptability, and Machiavellianism. As the BSF was similarly designed to index the *frequency with which one engages in bullshitting across various contexts in everyday life*, I felt a factor-level comparison would constitute a necessary and valuable test of whether the BSF truly measures a distinct construct. To that end, items from both scales were expected to best fit the data by loading onto separate factors.

Additionally, as bullshitting has been defined by some as "fall[ing] just short of lying" (Frankfurt, 1986; Stokke & Fallis, 2017), not only should scores on both the BSF and LiES scales be negatively associated with a measure of *lie acceptability,* but the strength of the association between BSF scores and lie acceptability should be significantly smaller than that of LiES scores. Though Study 2 showed that overall and persuasive bullshitting frequency was negatively associated with cognitive ability, lying has been found to be unrelated to cognitive

ability (Wright, Berry, Bird, 2012, 2013). Additionally, Frankfurt (1986, p. 16) has claimed that engagement in bullshitting involves processes that are "less deliberative and analytical" while Mears (2002) added that it is motivated by concerns related to self-image. Though research examining lying (and bullshitting) and analytic thinking is sparse, recent studies have found inconsistent associations with lying and self-esteem in adults (Harman, Hansen, Cochran, & Lindsey, 2005; Wright, White, & Obst, 2016). Therefore, I will also compare the associations between scores on the BSF and LiES with measures of cognitive ability, self-regard, and open-minded thinking to determine if the two scales can be further differentiated along these theoretical and empirical dimensions.

2.5.1 Method

2.5.1.1 Participants

Three hundred fifty-one adult participants from the United States and Canada were recruited from Amazon's Mechanical Turk participant pool using CloudResearch (Litman et al., 2016). My goal was to recruit approximately 300 participants, which would be sufficient for the factor analysis as well as provide ample power to detect an effect of r = .20 at an alpha = .05 (g*power; Faul, Erdfelder, Buchner, & Lang, 2009). Following the same exclusion protocols from Study 1, data from 34 participants were removed for violating Google reCAPTCHA v3 bot detection protocols and attention checks, leaving 317 participants for the final analysis (179 male, 136 female, 1 intersex, 1 prefer not to answer, M_{age} = 37.74, SD_{age} = 11.72, Bachelor's degree or higher = 54.8%).

2.5.1.2 Procedure

I preregistered my plan for Study 2 **(**https://osf.io/ak2vu). Participation procedures and restrictions were identical to those reported in Study 1. All participants completed the Bullshitting Frequency Scale (BSF) and the following measures, presented in a randomized order. Participants were paid $2.25 USD for their time.

2.5.1.3 Materials

Participants completed the 14-item *Lying in Everyday Situations* scale (LiES; Hart et al., 2019) to measure their propensity and motivations to lie in everyday situations, rating items on a 7-point "strongly disagree / strongly agree" scale. The LiES measures everyday lying across two factors: 1) *relational lying*, which reflects lying to avoid relational conflicts, and; 2) *antisocial lying*, involving lying that is intended to be harmful or vindictive. Participants also completed the 8-item *Revised Lie Acceptability Scale* (RLAS; Oliveira & Levine, 2008), which measures attitudes regarding how morally acceptable it is to lie to others, using a similar 7-point scale.

Cognitive ability was measured using the same materials as Study 1. Open-minded cognition and self-regard were measured using the Open-Minded Cognition Scale (Price et al., 2015), and the Core Self-Evaluations Scale (CSES; Judge et al., 2003).

2.5.2 Results

I first conducted a series of confirmatory factor analyses to both confirm the BSF's factor structure as well as ensure that the BSF and LiES scales were distinguishable at the factor level. I next calculated descriptive statistics as well as bivariate and partial Pearson's *r*-values for all

study variables (Table 7) and compared correlation coefficients using Fishers *r*-to–*z* transformations. I will first report the factor analyses, followed by the bivariate results.

2.5.2.1 Exploratory factor analysis

To test for possible cross-loading of items from both scales, I first conducted an exploratory factor analysis (EFA) using Principal Axis Factoring, entering all BSF and LiES items simultaneously. This method was chosen as I wanted to allow for the chance of significant cross-loadings between items, including the possibility that BSF items might load entirely on LiES scale factors (which would suggest either that the BSF measures lying rather than bullshitting, or that bullshitting and lying are psychometrically indistinguishable). Pattern matrix loadings after rotation are listed in Table 6. The sample yielded a Kaiser-Meyer-Olkin (KMO) score of .93, indicating "marvellous" sampling adequacy (Hutcheson & Sofroniou, 1999). Rotated pattern matrix scores loaded cleanly on four unique factors with eigenvalues of 10.86, 3.67, 2.23, and 1.12 (rotated component loadings of 8.35, 7.01, 7.07, and 4.97 respectively). Examination of the rotated component plot and scree plot justified this four-factor solution as best representing the data (Zwick & Velicer, 1982). Items loading on factors 1 and 4 corresponded with the *persuasive* and *evasive* subscales of the BSF while Factors 2 and 3 corresponded with those of the *relational* and *antisocial* subscales of the LiES scale (Table 6).

Table 6

Pattern matrix factor loadings after rotation for combined BSF and LiES scales

		1	2	3	4
1	BSF - When I want others to see me as more intelligent or knowledgeable.	.869			
2	BSF - When I want to impress the people I'm talking to.	.815			
3	BSF - When I'm trying to fit in better or be more accepted by the person or people I'm interacting with.	.744			
4	BSF - When I want the thing(s) I'm talking about to sound more interesting or exciting.	.730			
5	BSF - When I know it will help me achieve a goal.	.715			
6	BSF - When I know it will be easy to get away with it.	.679			
7	BSF - By pretending to know more about a topic than I actually do.	.629			
8	BSF - When I want to contribute to a conversation or discussion even though I'm not well-informed on the topic.	.516			
9	LIES - I tell lies in order to hurt, annoy, or upset others.		.868		
10	LIES - I lie for revenge.		.848		
11	LIES - I lie to people because it is amusing.		.844		
12	LIES - I lie in order to take people down.		.843		
13	LIES - I use lies to attack people I don’t like.		.839		
14	LIES - I lie because it is exciting.		.790		
15	LIES - I lie in order to punish people.		.782		
16	LIES - I lie to stay out of arguments with people.			.850	
17	LIES - I tell lies so I will not have confrontations with people.			.847	
18	LIES - I lie in order to hide shameful things about myself.			.774	
19	LIES - I lie in order to escape conflicts or disagreements with other people.			.746	
20	LIES - I lie to hide the bad things I’ve done.			.739	
21	LIES - I tell lies in order to spare another’s feelings.			.728	
22	LIES - I lie in order to be friendly and cordial with others.			.610	
23	BSF - When I don't want to tell someone what I really think.				.736
24	BSF - When being fully honest would be harmful or embarrassing to me or someone else.				.689
25	BSF - When someone asks me something that I want to avoid giving a direct answer to.				.659
26	BSF - When I need to fake/bluff my way out of a conversation or situation that I don't want to be in.				.613

Extraction Method: Principal Axis Factoring. *Rotation Method*: Oblimin with Kaiser Normalization. Rotation converged in 6 iterations. Sorted by size.

2.5.2.2 Confirmatory factor analyses

I next conducted a confirmatory factor analysis (CFA) using JASP (v0.11.1.0) on the BSF and confirmed that a two-factor model was a better fit for the data than a one-factor model (one factor: $\chi^2(54) = 269.49$, $p < .01$; CFI = .90; TLI = .88; RMSEA = .11; two factor: $\chi^2(53) = 132.29$, $p < .01$; CFI = .96; TLI = .95; RMSEA = .07). I next tested four combined BSF-LiES models: 1) a one-factor model combining all items from the BSF and LiES; 2) a two-factor model examining the BSF and LiES as separate factors; 3) a three-factor model examining the BSF as one factor and the LiES subscales as separate factors, and; 4) a four-factor model examining the BSF and LiES subscales all as separate factors. Results revealed that a 4-factor solution was the best fit for the data ($\chi^2(293) = 851.56$, $p < .01$; CFI = .91; TLI = .90; RMSEA = .08), suggesting that the BSF and LiES scales are distinct.

2.5.2.3 Bivariate correlations

Scores on the overall BSF and LiES scales (Table 7) were significantly and positively correlated, $r(315) = .65$, $p < .001$. Correlations for *persuasive*, $r(315) = .63$, $p < .001$, and *evasive* bullshitting, $r(315) = .54$, $p < .001$, with LiES scores followed the same pattern. Scores for both the BSF, $r(315) = .40$, $p < .001$, and the LiES scale, $r(315) = .54$, $p < .001$ were significantly and positively associated with RLAS scores (i.e., "lie acceptability"). These correlations were compared using Fisher's *r*-to-*z* transformation and found to be significantly different, $|z| = 2.26$, $p = .01$.

Scores for the BSF, $r(315) = -.22$, $p < .001$, and the LiES, $r(315) = -.37$, $p < .001$, were also significantly and negatively related to self-regard (CSES), and the difference between these correlations was statistically significant, $|z| = 2.06$, $p = .02$. BSF scores, $r(315) = -.26$, $p < .001$,

and LiES scores, $r(315) = -.32$, $p < .001$, were also significantly and negatively related to open-minded thinking (OMC), though the difference in correlations was not significant, $|z| = 0.82$, $p = .21$. Finally, both BSF, $r(315) = -.14$, $p = .02$, and LiES scores, $r(315) = -.11$, $p = .051$, were negatively related to cognitive ability, though the latter association was not significant nor was the difference between the correlations, $|z| = 0.38$, $p = .35$.

2.5.2.4 Partial correlations

I next examined the partial correlations of all variables with each of the BSF subscales (i.e., persuasive and evasive) controlling for the other subscale. *Persuasive bullshitting* scores (controlling for evasive) were significantly and positively related to overall lying, $r(314) = .42$, $p < .001$, relational lying, $r(314) = .32$, $p < .001$, antisocial lying, $r(314) = .35$, $p < .001$, and lie acceptability, $r(314) = .20$, $p < .001$. Persuasive bullshitting was also negatively related to open-minded cognition, $r(314) = -.16$, $p < .001$, and cognitive ability, $r(314) = -.12$, $p = .02$. *Evasive bullshitting* scores (controlling for persuasive) were significantly related only to overall lying, $r(314) = .21$, $p < .001$, relational lying, $r(314) = .28$, $p < .001$, and lie acceptability, $r(314) = .15$, $p < .001$. Fisher's *r*-to-*z* transformations revealed that *evasive bullshitting* was significantly different from *persuasive bullshitting* on overall lying, $|z| = 2.92$, $p < .01$, and antisocial lying, $|z| = 4.43$, $p < .01$.

Table 7

Descriptive and intercorrelational data for BSF and LiES scales with other study variables

		M	*SD*	α	*Bivariate*						*Partial*	
					BSF	BSFp	BSFe	LiES	LiES-R	LiES-A	BSFp[a]	BSFe[b]
1	BSF (overall)	2.57	0.76	*.92*	-							
2	Persuasive bullshitting (BSFp)	2.48	0.81	*.92*	.91**	-						
3	Evasive bullshitting (BSFe)	2.65	0.85	*.82*	.92**	.67**	-					
4	LiES (overall)	2.85	1.04	*.91*	.64**	.63**	.54**	-			.42**	.21**
5	Relational lying (LiES-R)	3.79	1.44	*.92*	.61**	.56**	.55**	.87**	-		.32**	.28**
6	Anti-social lying (LiES-A)	1.90	1.10	*.94*	.41**	.45**	.31**	.76**	.34**	-	.35**	.01
7	Lie acceptability	3.44	1.18	*.89*	.40**	.37**	.35**	.54**	.56**	.28**	.20**	.15**
8	Open-minded cognition	5.04	1.11	*.82*	-.25**	-.25**	-.20**	-.32**	-.21**	-.32**	-.16**	-.05
9	Self-regard	3.46	0.77	*.91*	-.22**	-.20**	-.20**	-.37**	-.43**	-.15**	-.08	-.10
10	Cognitive ability	7.81	1.55	*.73*	-.13*	-.15**	-.09	-.11	.06	-.28**	-.12*	.02

Note: N = 317. BSF = Bullshitting Frequency Scale; LiES = Lying in Everyday Situations; BSFp[a] = Persuasive bullshitting, controlling for evasive; BSFeb = Evasive bullshitting, controlling for persuasive.

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

2.5.3 Discussion

Results from Study 4 achieved two goals. First, confirmatory factor analyses provided further support for the two-factor structure of the BSF. Second, separate CFA models revealed that items from the BSF and the LiES scales clearly load on separate factors, indicating that they measure distinguishable constructs. Furthermore, the association of BSF scores and LiES scores with a measure of lie acceptability were significantly different. That is, more frequent bullshitters found lying less morally acceptable than more frequent liars. This supports the claim that bullshitters stop "just short of lying," and that liars are significantly more willing than bullshitters to intentionally convince people of falsehoods. Additionally, the *persuasive* and *evasive* bullshitting subscales were found to differ in their association with antisocial lying, with persuasive bullshitting (controlling for evasive) positively related while evasive (controlling for

persuasive) was unrelated. This supports the notion that evasive bullshitting is often motivated by a desire to avoid social harm (for self or others).

With respect to the remaining variables, liars reported experiencing significantly lower self-regard than bullshitters. Indeed, when examining the subscales for both measures, the effect size of self-regard was approximately twice as large for liars than bullshitters. Though I made no specific prediction about the associations with self-regard (other than the expectation that both would be negative), the finding that liars have significantly lower self-regard than bullshitters fits with the idea that individuals with lower self-esteem may be more strongly motivated to engage in more extreme forms of social manipulation. Lastly, bullshitting (overall and persuasive) and lying were both negatively related to open-minded thinking and cognitive ability. Though this latter association was not significant for overall or relational lying, it is interesting to note that it was significant for antisocial lying.

Overall, these results support the BSF as a measure of two types of "everyday bullshitting" that are distinguishable from everyday lying, and represents a valid tool for differentiating between these constructs. That said, there is clearly substantial overlap between "bullshitters" and "liars" as is evident in both the correlation between the BSF and LiES scales and the consistent (directional) relation with other constructs. This is consistent with the conceptual overlap discussed in the introduction and will represent an important consideration for future research investigating these constructs.

2.6 Study 5 – Persuasive and evasive bullshitting

My goals in Study 5 were twofold. First, I wanted to test four reworded BSF items that I felt might enhance the content validity and factor structure of each subscale. Additionally, I sought to further distinguish *persuasive* and *evasive bullshitting.* My previous studies showed that persuasive and evasive bullshitting significantly differ in their associations with measures of honesty, antisocial lying, open-minded cognition, and cognitive ability. However, an important test of the extent to which these two types of bullshitting frequency are discriminable would be to examine the associations of each subscale to performance on tasks that are more representative of actual "bullshitting behaviour." To that end, in addition to the BSF, I presented participants in Study 5 with two types of tasks.

The first of these comprised a political overclaiming questionnaire (OCQ) and a proposition-based overclaiming task (Dunlop, et al 2019), both of which measure exaggerated claims of political knowledge. I found some evidence in Study 1 that persuasive bullshitting was positively related to an overclaiming task while evasive was not, so I sought to examine if these results would generalize to a different knowledge domain. Critically, the proposition-based task goes beyond simply claiming one has familiarity with a topic, requiring them, if they claim knowledge, to articulate it in a provided text box. Thus, in addition to extending the overclaiming results to a novel domain, this task raises the proverbial stakes for respondents in requiring them to engage in the very act of supporting their beliefs.

The second test involves a novel social decision task which presents participants with various "real-world" scenarios and asks them the likelihood that they would tell the truth, lie, or be evasive when confronted with those types of situations in real life. If the persuasive and

evasive subscales of the BSF truly represent distinguishable constructs, then scores for each should significantly diverge in terms of their associations with responses on these two tasks. Specifically, *persuasive* (but not evasive) bullshitting should be positively related to the overclaiming tasks, given that exaggerated claims of one's knowledge or expertise are, by definition, examples of persuasive bullshitting. Likewise, *evasive* (but not persuasive) bullshitting should be positively related to evasive responses on the social decision task given that evasive bullshitting, by definition, involves avoiding potentially harmful (to self or others) direct responses by substituting evasive, non-relevant truths.

2.6.1 Method

2.6.1.1 Participants and procedure

Preregistration for Study 5 can be found at https://osf.io/xat6q. Three hundred ninety-eight adult participants from the United States and Canada were recruited from Amazon's Mechanical Turk participant pool using CloudResearch (Litman et al., 2016). Participation restrictions were identical to those of the previous studies. All participants completed the Bullshitting Frequency Scale (BSF) and the following measures, presented in a randomized order. Participants were paid $3.00 USD for their time.

2.6.1.2 Materials

Participants completed a politically-themed version of the *Overclaiming Questionnaire* (OCQ-P; Dunlop et al., 2019) which contained 12 targets and six foils relating to the topic of politics and political concepts. Participants rated each item using a 3-point response scale from 0 ("I have never heard of this item...) to 2 ("I could talk intelligently to others about this

item/concept"). An overclaiming score was calculated by subtracting proportion of hits (i.e., number of real items a person claimed knowledge of) from proportion of false alarms (i.e., the number of fake/foil items a person claimed knowledge of). An additional political overclaiming task was also used, adapted from Study 3 of Dunlop et al. (2019). Participants were presented with 10 propositions, 3 of which were foils, which asked participants to rate whether they agreed, disagreed, or did not know enough about the proposition to have an opinion. For items rated agree or disagree, a text box appeared which allowed participants to write a short summary of their reasons for that response. Overclaiming scores for the propositions were calculated identically to those for the OCQ. I also pre-registered "total word count" for each text response as a possible exploratory variable to consider but concluded that a third measure of overclaiming was unnecessary. This data is available in the raw data files (https://osf.io/dh6vj/) but will not be discussed further.

Participants also completed a novel *social decision task* designed specifically for this study (see appendix). Each participant read four individually-presented vignettes describing common social interactions. Participants were asked to read each vignette and then evaluate three possible responses (i.e., truth, lie, evasive) according to how likely they would be to give each response (from 1 "definitely not" to 5 "yes, definitely") were they to encounter such a situation in real life. Mean scores were calculated for each response type. Finally, cognitive ability was measured following the procedures from Studies 1 and 2.

2.6.2 Results

Per the pre-registered exclusion guidelines, data from 98 participants were removed for failing Google reCAPTCHA v3 bot detection protocols, attention checks, or providing highly

unusual or irrelevant responses to open-ended numeracy items (e.g., responding with "GOOD AND USE FULL SURVEY," or copying/pasting the question as the answer). This left us with data for 300 participants for the final analysis (62% male, 36.7% female, 1.3% intersex or prefer not to answer, M_{age} = 36.28, SD_{age} = 10.89, Bachelor's degree or higher = 57.3%). Bivariate and partial correlations, as well as descriptive statistics, for all study variables are listed in Table 9.

2.6.2.1 Confirmatory analyses of new BSF items

Principal axis factoring results showed that three of the new/reworded items loaded higher on their respective factors than 3 previous items, so the older items were discarded. The new version of the scale resulted in an eigenvalue of 6.71 for the *persuasive* factor (accounting for 55.93% of the variance) while the *evasive* factor improved to an eigenvalue of 1.44 (accounting for 12.0% of the variance, nearly double the previous version of the scale). The factor correlation for this updated version of the BSF was r = .59, compared to r = .76 for the previous version from Study 1 (Table 8).

I next sought to confirm that a two-factor structure remained the better fit for the data by conducting a confirmatory factor analysis (CFA) using JASP (v0.11.1.0). Results confirmed that the two-factor model ($\chi^2(53) = 136.23$, $p < .01$; CFI = .96; TLI = .95; RMSEA = .07) remains a better fit to the data compared to a one-factor model ($\chi^2(54) = 356.33$, $p < .01$; CFI = .86; TLI = .83; RMSEA = .14).

Table 8

Pattern matrix factor loadings after rotation for each scale item

		1	2	*M*	*SD*
1	When I want to impress the people I'm talking to.	.875		2.38	0.99
2	When I want others to see me as more intelligent or knowledgeable.	.860		2.41	1.09
3	When I want to contribute to a conversation or discussion even though I'm not well-informed on the topic.	.836		2.21	1.02
4	By pretending to know more about a topic than I actually do.	.810		2.22	1.01
5	When I'm trying to fit in better or be more accepted by the person or people I'm interacting with.	.779		2.42	1.09
6	When I know it will be easy to get away with it.	.667		2.25	1.11
7	When I want the thing(s) I'm talking about to sound more interesting or exciting.	.633		2.64	1.07
8	When I'm trying to persuade someone to change their mind or agree with what I'm saying.	.630		2.40	1.00
9	When being fully honest would be harmful or embarrassing to me or someone else.		.871	2.93	1.06
10	When a direct answer would hurt another person's feelings.		.855	2.74	1.06
11	When a direct answer might get me in trouble.		.572	2.81	1.04
12	When I don't want to tell someone what I really think.		.542	3.07	1.04

Extraction Method: Principal Axis Factoring. Rotation Method: Oblimin with Kaiser Normalization. Rotation converged in 5 iterations. Factor correlation, $r = .59$

2.6.2.2 Bivariate correlations

As my hypotheses were concerned specifically with distinguishing the BSF subscales, my discussion of the results will focus on the subscales only, though data for the overall BSF scale is listed in Table 9. *Persuasive* and *evasive bullshitting* were correlated at the bivariate level $r(300) = .62$, $p < .001$. Additionally, *persuasive bullshitting* was significantly and positively related to overclaiming, $r(300) = .18$, $p < .001$, proposition overclaiming, $r(300) = .15$, $p = .01$, and lie responses on the social decision task, $r(300) = .40$, $p < .001$, and significantly and negatively related to cognitive ability, $r(300) = -.31$, $p < .001$. *Evasive bullshitting* was significantly and

positively related to lie responses, $r(300) = .39$, $p < .001$, and evasive responses, $r(300) = .12$, $p = .04$, on the social decision task and significantly and negatively related to cognitive ability, $r(300) = -.16$, $p < .001$.

2.6.2.3 Partial correlations

Turning to the partial correlations, *persuasive bullshitting* (controlling for evasive) was significantly and positively related to overclaiming, $r(297) = .26$, $p < .001$, proposition overclaiming, $r(297) = .15$, $p < .001$, truthful responses, $r(297) = .15$, $p = .01$, and lie responses, $r(297) = .22$, $p < .001$. *Persuasive bullshitting* was also significantly and negatively related to evasive responses, $r(297) = -.13$, $p = .03$, and cognitive ability, $r(297) = -.27$, $p < .001$. *Evasive bullshitting* (controlling for persuasive) was significantly and positively related to lie responses, $r(297) = .19$, $p < .001$, and evasive responses, $r(297) = .17$, $p < .001$, and significantly and negatively related to overclaiming, $r(297) = -.19$, $p < .001$, and truthful responses, $r(297) = -.17$, $p < .001$.

Table 9

Descriptive and correlational data for all study variables with BSF, BSFp, and BSFe

		M	*SD*	α	*Bivariate*			*Partial*	
					BSF	BSFp	BSFe	BSFp[a]	BSFe[b]
1	BSF (overall)	2.63	0.78	*.92*	-			-	-
2	Persuasive bullshitting (BSFp)	2.36	0.85	*.93*	.90**	-		-	-
3	Evasive bullshitting (BSFe)	2.89	0.87	*.85*	.90**	.62**	-	-	-
4	Overclaiming	-0.66	0.46	*.88*	.08	.18**	-.03	.26**	-.19**
5	Proposition overclaiming	-0.46	0.23	*.64*	.10	.15*	.05	.15**	-.06
6	Social response - truth	2.08	0.85	*.73*	-.03	.05	-.10	.15*	-.17**
7	Social response - lie	3.02	0.90	*.75*	.44**	.40**	.39**	.22**	.19**
8	Social response - evasive	3.69	0.76	*.63*	.05	-.03	.12*	-.13*	.17**
9	Cognitive ability	7.95	1.47	*.75*	-.26**	-.31**	-.16**	-.27**	.04

Note: N = 300. BSF = Bullshitting Frequency Scale; BSFp[a] = Persuasive bullshitting, controlling for evasive; BSFe[b] = Evasive bullshitting, controlling for persuasive. ***p* < .01; **p* < .05

2.6.3 Discussion

Consistent with my predictions, the results from Study 5 provide strong evidence that the *persuasive* and *evasive bullshitting* subscales of the BSF measure clearly distinguishable constructs. Indeed, each subscale (when controlling for the other) was positively associated with performance on conceptually congruent tasks and negatively associated with performance on conceptually incongruent tasks. That is, the three measures of overclaiming, truthful responses, and evasive responses provide a complete dissociation, in that the correlations for these variables with each bullshitting subscale were significant but in opposite directions.

Specifically, *persuasive bullshitting* scores were positively associated with performance on two tasks measuring the propensity to exaggerate (or otherwise positively misrepresent) one's knowledge of political topics which, by definition, is a form persuasive bullshitting. Additionally, *persuasive bullshitting* scores were more likely to be associated with choosing both

direct/truthful and lie responses and negatively related to the likelihood of responding evasively when confronted with those situations. Given the negative associations between persuasive bullshitting, open-minded cognition, and cognitive ability, as well as the positive association with relational lying (see Study 4), this suggests that persuasive bullshitters may lack sufficient analytic and/or theory of mind processes to quickly formulate effective evasive responses, so they instead opt for less cognitively effortful or complex responses (i.e., simple truths or white lies).

Likewise, *evasive bullshitting* was negatively related to two measures of overclaiming one's political knowledge, suggesting that they are perhaps less concerned with (or less motivated by) positively misrepresenting their personal qualities or ideas relative to high persuasive bullshitters. Additionally, *evasive bullshitting* was positively related to responding evasively (and by lying) in precarious social situations which, by definition, reflects evasive bullshitting. Interestingly, *evasive bullshitters* were also less likely to choose direct, truthful responses to socially precarious inquiries. However, this does not necessarily reflect dishonesty. Given both the definition of evasive bullshitting provided in the introduction, as well as the positive association with relational lying found in Study 4, this negative association with choosing the direct/truthful response likely simply reflects evasive bullshitters strong desire to avoid responding in a way that they perceive may lead to negative social costs (e.g., hurt feelings, embarrassment, etc.).

2.7 General Discussion

In the present investigation, I created the Bullshitting Frequency Scale (BSF), a valid measure which captures the self-reported propensity with which individuals engage in two types

of "everyday bullshitting," *persuasive* and *evasive*, and conducted an initial investigation into the cognitive and dispositional individual differences associated with the propensity to engage in this behaviour. Overall, bullshitting can be understood as an instrumental and performative communication strategy employed to either: 1) impress, persuade, or fit in with others by exaggerating one's knowledge, attitudes, skills, or competence (i.e., *persuasive bullshitting),* and/or; 2) attempts to evade or altogether avoid responding to inquiries where direct answers might result in negative social costs (i.e., *evasive bullshitting*).

This propensity was found to be negatively related not only to self-report measures of honesty, sincerity, open-mindedness, self-worth, and cognitive ability, but also positively related to actual behaviour as measured by performance on tasks that are (arguably) "bullshit congruent," specifically those that measure exaggerated claims of the depth of one's general and political knowledge and responses to precarious social situations. Additionally, when compared to a valid, reliable measure of everyday lying (i.e., LiES; Hart et al., 2019), the BSF demonstrated a unique factor structure which was also differentially related to perceived moral acceptability of lying and to self-regard. Overall, the results presented here highlight important cognitive and dispositional factors related to the propensity to engage in two types of bullshitting and suggest that such behaviour can be reliably measured using a self-report scale. It should be noted, though, that more non-Western and/or non-English-speaking cultures may have different conceptions of these types of behaviours, how they are interpreted, and what they represent, and future research would greatly benefit from exploring bullshitting cross-culturally (see Giles, Rothermich, & Pell, 2019).

2.7.1 Bullshitting vs lying

An important question to ask is whether bullshitting is distinct enough from lying to be theoretically interesting and worthy of empirical pursuit. In other words, is the construct of bullshitting (and its measurement with the BSF) simply "old wine in a new bottle?" As I have shown here, bullshitting and lying are related but distinct, the motivations behind bullshitting and lying are different, and "prolific bullshitters" differ from "prolific liars" in identifiable, measurable ways. Furthermore, I argue that the BSF represents another step forward in the empirical examination of bullshitting as a meaningful construct.

Williams (2002) gives a standard definition of lying as knowingly and intentionally making statements believed to be false. Bullshitting, on the other hand, amounts to arguably less severe distortions meant to impress, persuade, or evade which have, at most, a "loose concern for the truth" (Frankfurt, 1986; Mears, 2002; Meibauer, 2018). That is, the veracity of what is said is arguably less important to the bullshitter whereas it is crucially important to the liar. A key distinction is that the liar's intent is to deceive with falsehoods whereas the bullshitter's intent is to foster positive impressions (or avoid negative ones) using tactics such as exaggerations, embellishments, and evasions (Carson, 2016; Mears, 2002). The results of Study 2 provide support for this idea, in that the Bullshitting Frequency Scale (BSF) and the Lies in Everyday Situations scale (LiES) were shown to be factorially distinct (i.e., they measure different constructs). Additionally, liars were found to have significantly lower self-regard while bullshitters were found to be significantly less likely to view lying as a morally acceptable behaviour.

Consideration of these distinctions has largely been absent from prior deception research which has instead utilized broad descriptions that conflate the definitions of lying and bullshitting that philosophers, linguists, and some psychologists (e.g., Pennycook et al., 2015; Petrocelli, 2018) have earnestly attempted to distinguish. For instance, in their classic study on everyday lying, DePaulo et al., (1996), recorded all misleading statements from participants, "no matter how big or small" (p. 981), as instances of lying and rated them on a scale from "trivial and unimportant" to "serious and very important." Furthermore, subtle evasions, simple exaggerations, and "outright falsehoods" were all coded as "lies" in the data analyses. Subsequent research on lying has relied on similar paradigms (e.g. Feldman et al., 2002; Serota et al., 2009). Under such conditions, lies constitute anything from trivial exaggerations about unimportant topics (e.g., bragging about one's cooking skills) to serious, outright falsehoods about very important topics (e.g., a murder suspect lying to police).

This poses potential issues for individual differences research in that the cognitive and psychological factors (as well as the consequences) associated with the arguably less serious act of bullshitting (Mears, 2002) are likely to differ from the darker attributes related to more duplicitous, pathological lying (Jones & Paulhus, 2017). The BSF will help to advance knowledge in this area by allowing researchers to examine these constructs and their associated individual differences separately.

2.7.2 Being honest about bullshitting

It is also important to consider the extent to which one can be confident that participants responded honestly to items on the BSF and were not, in fact, bullshitting the researchers. Here, I have asked participants to respond honestly about the frequency with which they engage in

somewhat less-than-honest behaviour. The irony of such an exercise notwithstanding, I feel confident that data collected using the BSF accurately reflect, at least to a significant degree, real-world behaviour. Prior research has shown that, consistent with the present results, only a small proportion of people report that they frequently engage in dishonest behaviour (Halevy, Shalvi, & Verschuere, 2014; Serota & Levine, 2014). Other work has shown that self-report and other-report data for measures of honesty and social desirability are strongly positively correlated (de Vries, Lee, & Asthon, 2008; de Vries et al., 2014; Halevy et al., 2014). Additionally, self-reports of social desirability and lying/dishonesty significantly and positively correlate with actual cheating (Halevy et al., 2014; Zettler et al., 2015). Indeed, a similar pattern was found in the present data, where persuasive bullshitting frequency was positively associated with performance on behavioural tasks (i.e., overclaiming). Given this sizeable body of work confirming a consistent agreement in self-report, other-report, and behavioural measures using scales that are highly similar, if not identical, to measures I used in this study, I feel confident that this data represent overall honest and accurate responding patterns from participants.

2.7.3 Conclusion

The present results support the Bullshitting Frequency Scale as a valid, reliable measure of bullshitting and illuminate important cognitive and dispositional factors associated with the propensity to bullshit others. The results here represent an important step forward in the psychological study of our receptivity to falling for and propensity for producing vacuous-yet-persuasive, evasive, or otherwise misleading statements (i.e., bullshit).

Chapter 3

Bullshitting frequency predicts receptivity to various types of misleading information

3.1 Introduction

"Most people are rather confident of their ability to recognize bullshit and to avoid being taken in by it." – Harry Frankfurt (2005|1986)

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.

3.1.1 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, I define bullshit, broadly, as information designed to impress, persuade, and/or otherwise mislead that is often constructed with an indifference for the truth or meaning (Carson, 2016; Frankfurt, 2005; Gligorić et al., 2020; Pennycook et al., 2015; Reisch, 2006). 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).

3.1.2 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.

3.1.3 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, I report here four studies focused on examining the associations between bullshitting frequency and bullshit receptivity. In Studies 6 and 7, I 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. I follow up these correlational results experimentally in Studies 8 and 9 by investigating potential mechanisms underlying this association. Data files for all studies are available here: https://osf.io/chpvm/.

3.2 Study 6

In Study 6, I 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.[2] Each bullshit task also includes a measure of receptivity to contextually-relevant non-bullshit (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).

3.2.1 Method

3.2.1.1 Participants

I recruited 261 adult participants from the United States and Canada from Amazon's Mechanical Turk participant pool using the crowdsourcing platform, Cloudresearch (Litman,

[2] It should be noted that Study 6 was conducted after Studies 7 through 9. However, we present it here first, as we feel that this presentation order provides helpful theoretical context to the reader for the remainder of the chapter.

Robinson, & Abberbock, 2016). To meet my goal of achieving at least .80 power to detect an effect of r = .20 at α = .05, an a priori power analysis indicated that I would need a sample of 191, which this 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 data for 219 participants to consider in the final analyses (127 male, 91 female, 1 prefer not to answer, M_{age} = 37.94, SD_{age} = 11.44, Bachelor's degree or higher = 62%). Participants were paid $2.00 USD for the roughly 15-minute study.

3.2.1.2 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):

3.2.1.3 Materials

Bullshitting Frequency

To assess the self-reported frequency with which a person utilizes two types of bullshitting in various contexts, I 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 5-point scale (from "not at all profound" to "very profound"), the profundity of 10 randomly generated, yet grammatically correct, sentences that were constructed from abstract pseudo-profound 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), I 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.

3.2.2 Results

Descriptive statistics and bivariate correlations with persuasive and evasive bullshitting frequency are listed in Table 10 (intercorrelations for all variables can be found in the supplementary materials). Multiple linear regression models were created to examine the extent

to which these variables could predict each type of bullshit receptivity (Table 11). I focus first on the correlations and then on the linear regression models.

3.2.2.1 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 the bullshit receptivity measures. To examine the ways in which persuasive and evasive bullshitting might be differentially related to the other variables, I 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 10

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

		M	*SD*	α	*Bivariate*		*Partials controlling for BSF subscales*	
					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. $^{**}p < .01$; $^{*}p < .05$

3.2.2.2 Linear regressions

To test my main question of whether propensity to engage in bullshitting predicts receptivity to various types of bullshit, I created three multiple linear regression models (Table 11), 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, I 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 the 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.[3]

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), $\beta = -.16$, $p = .02$, 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].

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

Table 11

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

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

3.2.3 Discussion

In Study 6, I 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, 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.

3.3 Study 7

In Study 7, I examine more deeply the extent to which bullshitting frequency is associated with receptivity to misleading information (i.e., bullshit). To that end, I limit my 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, I 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).

I also include a number of additional measures to allow for an exploration of 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 pseudo-profound bullshit receptivity.

3.3.1 Method

3.3.1.1 Participants

I 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 my goal of achieving at least .80 power to detect an effect of r = .20 at α = .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 data for 207 participants to consider in the final analyses (137 male, 69 female, 1 intersex, M_{age} = 36.75, SD_{age} = 11.18, Bachelor's degree or higher = 60.9%). Participants were paid $3.50 USD for the roughly 25-minute survey.

3.3.1.2 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). 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.

I also collected confidence ratings for each cognitive ability item, using a sliding scale from 0-100, from which I 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

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 (α = .91).

Metacognition (self-reported)

I 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 (α = .90; Littrell et al., 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 (α = .72; Krumrei-Mancuso & Rouse, 2016). In order to capture self-reported intellectual overconfidence, I reverse-scored the scale.

3.3.2 Results

3.3.2.1 Correlations

Correlations and descriptive statistics for key variables can be found in Table 12 (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, I 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 12

Descriptive and correlational data for BSF with all study variables

		M	*SD*	α	*Bivariate*			*Partial*	
					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	Profoundess 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; BSFp[a] = Persuasive bullshitting, controlling for evasive; BSFe[b] = Evasive bullshitting, controlling for persuasive.

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

3.3.2.2 Linear regression

I 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, I entered the two bullshitting frequency variables and receptivity to motivational quotes. In Step 2, I entered cognitive ability. Finally, in Step 3, I entered the cognitive reflection and metacognition variables (i.e., CRT, insight, and the two overconfidence variables). As self-reported reflection was not significantly associated with bullshit receptivity at the bivariate level, it was excluded from the regression model. The 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 13.

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 13

Multiple linear regressions for all study variables predicting receptivity to bullshit (BSR score)

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

3.3.3 Discussion

Study 7 delved deeper into the association between the propensity to engage in bullshitting and the propensity to fall for bullshit. As with Study 6, 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 the cognitive and metacognitive variables. For instance, persuasive bullshitting (controlling for evasive) was positively related to bullshit receptivity and the 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, the linear regression model supports the previous two studies' 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 even though 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 the 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 I 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.

3.4 Study 8

In Study 8, I focus on better understanding why receptivity to pseudo-profound bullshit is predicted by persuasive bullshitting frequency by examining a potential mechanism for this

relation. As noted in Study 7, the positive relation between persuasive bullshitting frequency and bullshit receptivity could not be explained by variables related to subjective and objective measures of various cognitive and metacognitive skills and abilities. This suggests that the extent to which a persuasive bullshitter is a reflective thinker and possesses clarity of thought does not inoculate him from the seductive nature of pseudo-profound bullshit. Given that persuasive bullshitting is negatively related to evaluations of self-worth and positively related to overclaiming (Littrell et al., 2020), it could be that persuasive bullshitters assign higher profoundness ratings to pseudo-profound items as a kind of low-cost self-enhancement or impression management strategy (Turpin et al., 2019). For example, more prolific persuasive bullshitters may view the ability to identify profoundness (or at least the *claim* that one has this ability) as socially beneficial (i.e., a signal to others that they are profound, deep thinkers). Under this view, rating pseudo-profound bullshit items as being "more profound" (relative to more mundane items) would itself be a strategic or instrumental form of bullshitting (Littrell et al., 2020; Reisch, 2006).

To test this hypothesis, I altered the wording of the bullshit receptivity scale (BSR) instructions for one group of participants in an attempt to discourage possible self-enhancement or impression management motivations. Using an approach based on previous work by Atir, Rosenzweig, and Dunning (2015), half of the participants received BSR instructions that included a warning that, despite how profound the items may seem, some of them were not actually profound and that participants should do their best to identify these pseudo-profound items when assigning profoundness ratings. I predicted that, if persuasive bullshitters indeed assign higher profoundness ratings to BSR items in order to appear more intellectually

perspicacious, then this warning should discourage such behaviour, thus decreasing the association between persuasive bullshitting and bullshit receptivity (relative to the ratings from the group who receive the standard BSR instructions).

3.4.1 Method

3.4.1.1 Participants

Based on a power analysis for achieving at least .80 power to detect an effect of r = .20 at α = .05 (g*power; Faul, Erdfelder, Buchner, & Lang, 2009), 212 adult participants from the United States and Canada were recruited from Amazon's Mechanical Turk participant pool via the CloudResearch crowdsourcing platform (Litman, Robinson, & Abberbock, 2016). Inclusion criteria were identical to those of Study 1. Data for 7 participants were removed for failing attention checks, and an additional 6 were removed for providing response patterns that were identified as potential "bots" by Google reCAPTCHA v3 bot detection protocols, leaving data for 199 participants to consider in the final analyses (115 male, 84 female, M_{age} = 36.20, SD_{age} = 10.61, Bachelor's degree or higher = 52.3%).

3.4.1.2 Procedure

I preregistered my hypotheses and methods (available at https://osf.io/4tmbk) on the Center for Open Science's "Open Science Framework" (OSF). All participants completed the following measures, with bullshitting frequency and cognitive ability items presented in a random order followed by the BSR. Random assignment of participants to the two BSR instruction conditions was carried out by Qualtrics' automated block randomization procedure.

3.4.1.3 Materials

Bullshitting Frequency and Bullshit Receptivity

Persuasive and evasive bullshitting frequency were again measured using the Bullshitting Frequency Scale (BSF). To measure receptivity to pseudo-profound bullshit, participants again completed the Bullshit Receptivity Scale (BSR; Pennycook et al., 2015). For the BSR, participants were randomly assigned to one of two instruction conditions.

Group 1 received the standard BSR instructions: *"For the following items, please rate how profound each statement is, on a scale of 1 (not at all profound) to 5 (very profound). The definition of profound is "showing great knowledge or insight; to be taken as deeply meaningful; of great and broadly inclusive significance."* The second group received alternate instructions which were composed of the standard instructions along with the following addendum: *"Several of the items were designed to sound profound but are not. Discriminating between profound and not profound statements can be difficult, so please do your best."* To reduce the possibility that participants might skip over these additional instructions without reading, the addendum was presented a second time on a screen by itself (for emphasis), directly after the primary instruction screen.

Cognitive ability

Following the procedure listed in Study 2, a mean cognitive ability score was calculated using combined scores from the General Risk and Numeracy Scale (Lipkus, Samsa, & Rimer, 2001) and the "Wordsum" vocabulary test (Thorndike, 1942; Malhotra, Krosnick, & Haertel, 2007).

3.4.2 Results

Descriptive statistics and Pearson's bivariate correlations for all variables are listed in Table 14. For the pseudo-profound (BSR) and profound/motivational (MQR) items, participants were randomly assigned to either the "Standard" (N = 102) or "Alternate" (N = 97) instruction conditions (see Materials). All participants (N = 199) completed the Bullshitting Frequency Scale (BSF) and the cognitive ability measures.

3.4.2.1 Bivariate correlations

In the "standard instruction" BSR condition, *persuasive bullshitting* was significantly and positively related to BSR-Standard, $r(100) = .36$, $p < .01$. In the "alternate instruction" BSR condition, *persuasive bullshitting* was significantly and positively related to BSR-Alternate, $r(95) = .58$, $p < .01$, and MQR-Alternate, $r(95) = .37$, $p < .01$. This pattern was also consistent for *evasive bullshitting* (BSR-Alternate, $r(95) = .47$, $p < .01$, and MQR-Alternate, $r(95) = .27$, $p < .01$).

3.4.2.2 Partial correlations

Persuasive bullshitting (controlling for evasive) was positively related to bullshit receptivity in the "standard instruction" condition, $r(99) = .36$, $p < .01$, and in the "alternate instruction condition," $r(94) = .47$, $p < .01$, and was positively related to receptivity to profound/motivational quotes in the "alternate instructions" condition, $r(94) = .26$, $p < .01$. In the "standard" condition, *evasive bullshitting* (controlling for persuasive) was negatively related to bullshit receptivity, $r(94) = -.18$, $p = .08$, and receptivity to profound/motivational quotes, $r(94) = .15$ $p = .13$, but neither of these associations was significant.

Table 14
Descriptive and correlational data for BSF with BSR and cognitive ability

		M	*SD*	α	*Bivariate*		*Partial*	
					BSFp	BSFe	BSFp[a]	BSFe[b]
1	Persuasive bullshitting	2.54	0.90	*.93*	-		-	-
2	Evasive bullshitting	2.74	0.91	*.85*	.76**	-	-	-
3	Bullshit receptivity (stnd)	27.59	9.54	*.92*	.36**	.17	.36**	-.18
4	Profoundness receptivity (stnd)	34.94	6.99	*.82*	.13	.20*	-.04	.15
5	Bullshit receptivity (alt)	25.58	10.21	*.93*	.58**	.47**	.39**	.07
6	Profoundness receptivity (alt)	33.79	8.60	*.88*	.37**	.27**	.26**	.00
7	Cognitive ability (overall)	7.27	1.97	*.83*	-.36**	-.25**	-.27**	.04

Note: N = 199 (overall), N = 102 (standard), N = 97 (alternate). BSFp[a] = Persuasive bullshitting, controlling for evasive; BSFe[b] = Evasive bullshitting, controlling for persuasive. $^{**}p < .01$; $^{*}p < .05$

3.4.2.3 Linear regression

I next created a multiple linear regression model predicting overall BSR scores from persuasive bullshitting (BSFp), evasive bullshitting (BSFe), profound/motivational quote receptivity (MQR), and cognitive ability. To test whether the instruction manipulation produced a measurable effect, I included a variable for BSR instruction condition as well as interaction terms for both persuasive (PersuasiveBS*condition) and evasive (EvasiveBS*condition) bullshitting. All continuous predictor variables were mean-centred prior to inclusion in the regression model. Beta coefficients (with 95% CI) and model fit statistics are listed in Table 15. Plots of interaction effects can be found in Figure 2.

Table 15

Multiple linear regressions for all study variables (mean-centered) predicting receptivity to bullshit (BSR score)

	b	*SE*	β	95% CI	
				Lower	Upper
(Constant)	27.42**	0.72			
Persuasive bullshitting (BSFp)	3.76**	1.25	.34**	.12	.57
Evasive bullshitting (BSFe)	-2.65*	1.18	-.24*	-.46	-.03
Profoundness receptivity (MQR)	.31**	0.07	.24**	.14	.35
Cognitive ability	-2.17**	0.28	-.43**	-.54	-.32
Condition	-1.36	1.04	-.07	-.17	.03
PersuasiveBS*condition	-0.62	1.78	-.04	-.26	.18
EvasiveBS*condition	3.77*	1.76	.23*	.02	.44
Adjusted R^2	.46				
F	25.52**				

Note: N = 199. CI = confidence intervals for standardized betas; $^{**}p < .01$; $^{*}p < .05$

Pseudo-profound bullshit receptivity (BSR) was significantly and positively predicted by persuasive bullshitting frequency (BSFp), $\beta = .34$, $p < .01$, 95% CI [.12, .57] and profound/motivational quote receptivity (MQR), $\beta = .24$, $p < .01$, 95% CI [.14, .35]. Bullshit receptivity was also significantly and negatively predicted by evasive bullshitting frequency (BSFe), $\beta = -.25$, $p = .03$, 95% CI [-.46, -.03], cognitive ability, $\beta = -.43$, $p < .01$, 95% CI [-.54, -.32]. For the interaction terms, only the EvasiveBS*condition term was a significant predictor, $\beta = .23$, $p < .01$, 95% CI [.02, .44]. To follow up on the interaction, I conducted simple slopes regression analyses for each instruction condition. Results revealed that, while persuasive bullshitting was a significant positive predictor in both conditions (Standard: $\beta = .37$, $p = .003$; Alternate: $\beta = .31$, $p = .014$), evasive bullshitting was only a significant negative predictor of bullshit receptivity in the standard instruction condition (Standard: $\beta = -.29$, $p = .016$; Alternate: $\beta = .09$, $p = .413$).

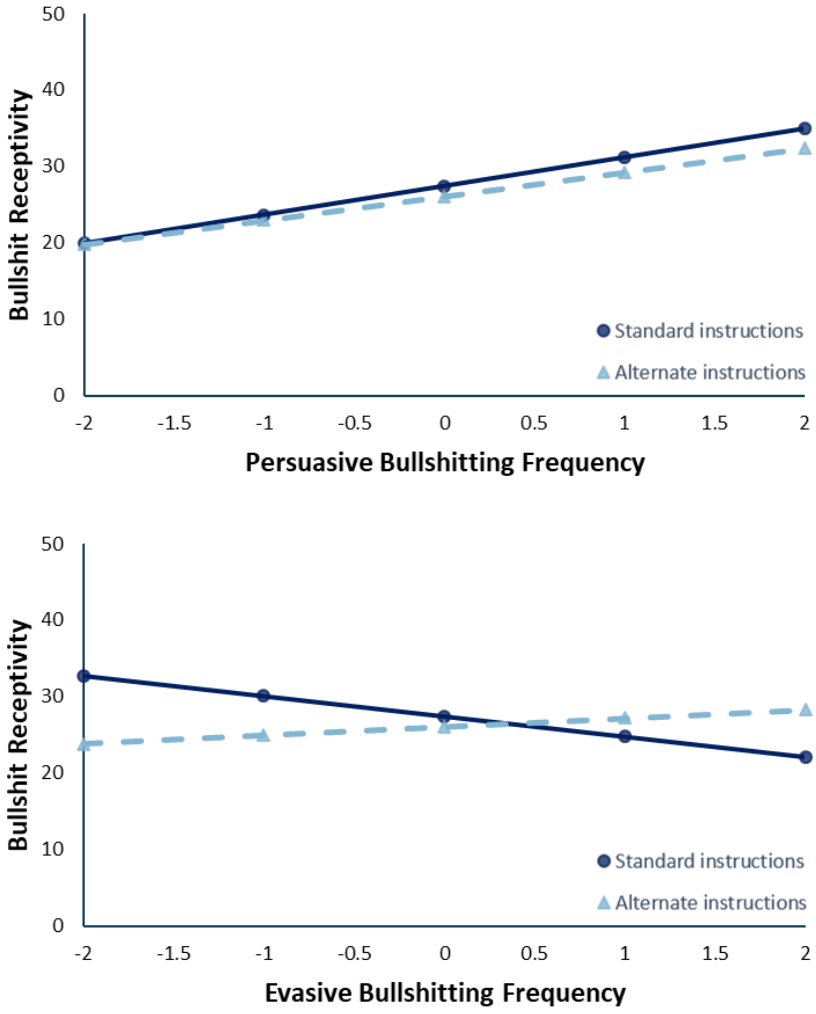

Figure 2. Plots of interaction effects for persuasive (top) and evasive (bottom) bullshitting predicting bullshit receptivity (BSR), controlling for all other variables. Dark solid lines represent standard instruction condition. Light dashed lines represent alternate instruction condition.

3.4.3 Discussion

In Study 8, I further examined the positive associations between persuasive bullshitting frequency and bullshit receptivity. The linear regression results again confirmed that *persuasive bullshitting* predicted bullshit receptivity, controlling for evasive bullshitting, profound/motivational quote receptivity and cognitive ability. The interaction between

persuasive bullshitting and condition, however, was not significant. A closer look at the regression models for each instruction condition revealed that persuasive bullshitting positively predicted BSR scores in both. That is, no matter the instruction, individuals with high persuasive bullshitting scores gave higher profoundness ratings to BSR items (Figure 2). This finding suggests that the relation between persuasive bullshitting and bullshit receptivity is not a product of an attempt at posturing by persuasive bullshitters. This result (seemingly) falsifies the hypothesis that assigning higher profoundness ratings to BSR items was, in essence, an instrumental form of bullshitting employed by persuasive bullshitters, and instead suggests that a more fundamental mechanism may be involved. It is also possible, of course, that the manipulation was ineffective.

The idea that the manipulation I used was ineffective, however, seems inconsistent with the fact that I found a significant interaction between the instruction condition and *evasive bullshitting* (Figure 2). That is, the negative association between bullshit receptivity and evasive bullshitting found in Studies 6 and 7 was replicated in Study 8 under the standard instructions (even when controlling for MQR and cognitive ability, which was not the case in Study 7) but this negative relation disappeared in the alternate instruction condition (indeed, the relation became slightly positive, though it was not significant). Given this result was not predicted, it is important to be cautious in placing too much stock in it.

3.5 Study 9

Studies 6 thru 8 have established a consistent, robust positive association between persuasive bullshitting frequency and bullshit receptivity, but the mechanism underlying this relation remains unclear. In Study 9, I test another 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, I 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.

3.5.1 Method

3.5.1.1 Participants

To ensure that I achieved at least .80 power to detect an effect of $r = .20$ at $\alpha = .05$ in both conditions, I recruited 454 adult participants from the United States and Canada from 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 a sample of 400 to consider in the final analyses. Participants were paid $0.75 USD for the roughly 10-minute study.

3.5.1.2 Procedure

My 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.

3.5.1.3 Materials

The materials used to measure bullshitting frequency and cognitive ability were identical to Study 3. 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.

3.5.2 Results

Here, I focus the results and discussion on the linear regression analyses of the experimental manipulation. However, descriptive statistics and intercorrelations for all variables can be found in Table 16. Following previous procedures, I 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 17). I included a variable for BSR instruction condition and interaction terms for both of the bullshitting frequency variables (see Figure 3). All predictor variables (excluding condition) were mean-centred.

Table 16
Descriptive and correlational data for bullshitting frequency with all study variables

		M	*SD*	α	*Bivariate*		*Partial*	
					BSFp	BSFe	BSFp[a]	BSFe[b]
1	Persuasive bullshitting (BSFp)	2.47	0.85	*.93*	-		-	-
2	Evasive bullshitting (BSFe)	2.63	0.84	*.81*	.70**	-	-	-
3	Bullshit receptivity ("sounds")	29.91	9.24	*.90*	.11	.01	.14	-.08
4	Profoundness recepitivy ("sounds")	33.99	7.04	*.82*	.10	.16*	-.01	.13
5	Bullshit receptivity ("is")	25.69	9.19	*.91*	.23**	-.01	.33**	-.25**
6	Profoundness recepitivy ("is")	32.78	7.79	*.85*	.17*	.07	.17*	-.07
7	Cognitive ability	7.81	1.36	*.83*	-.16**	-.07	-.28**	.12

Note: N = 400 (overall), N = 201 ("Sounds profound"), N = 199 ("Is profound"). **$p < .01$; *$p < .05$

3.5.2.1 Linear regressions

Neither persuasive, β = .09, p = .26, 95% CI [-.07, .25] nor evasive, β = -.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, β = .18, p = .03, 95% CI [.02, .35], though the EvasiveBS*condition interaction was not significant (Table 17). Following up on this finding, I created regression models for each instruction condition (Figure 3). This revealed that in the "is profound" condition, both persuasive, β = .37, p < .001, 95% CI [.21, .53], and evasive, β = -.31, p < .001, 95% CI [-.47, -.15], bullshitting significantly predicted bullshit receptivity (controlling for MQR and cognitive ability), matching the pattern found in the previous three studies for the standard instructions. However, neither persuasive, β = .12, p = .19, 95% CI [-.06, .30], nor evasive, β = -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 3) 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 17

Multiple linear regressions for all study variables (mean-centered) predicting receptivity to bullshit (BSR score)

	b	*SE*	*β*	95% CI	
				Lower	Upper
(Constant)	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	.18*	.02	.35
EvasiveBS*condition	-1.78	1.31	-.11	-.28	.05
Adjusted R^2	.33				
F	28.98**				

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

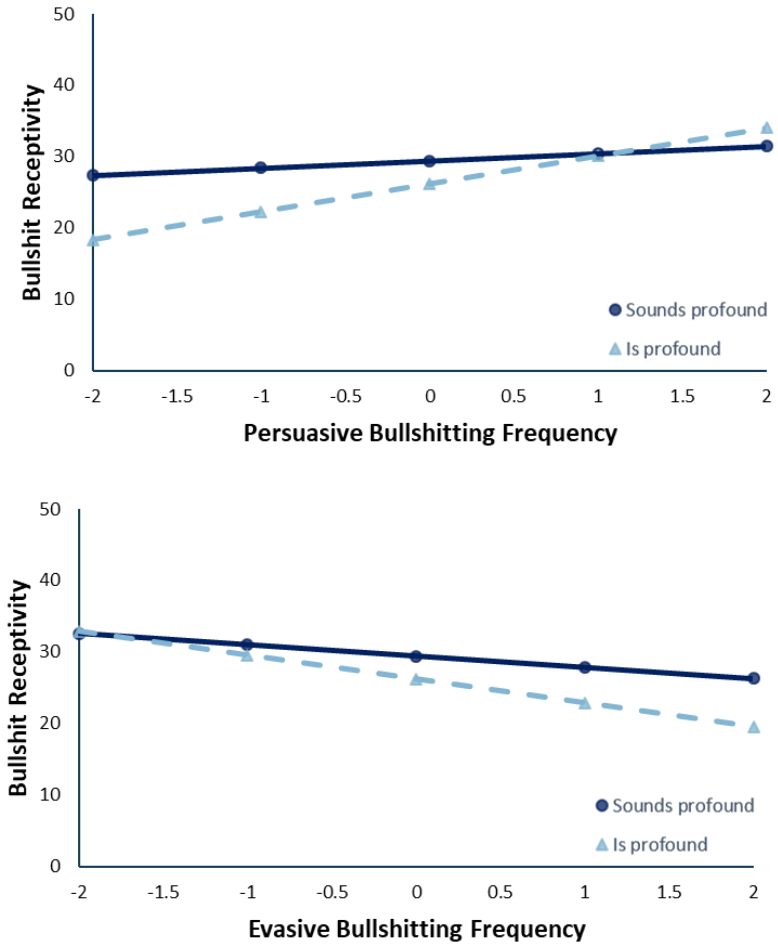

Figure 3. Plots of interaction effects for persuasive (top) and evasive (bottom) bullshitting predicting bullshit receptivity (BSR), controlling for all other variables. Dark solid lines represent "Sounds profound" instruction condition. Light dashed lines represent "Is profound." instruction condition.

3.5.3 Discussion

The goal for Study 9 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, the 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."

3.6 General discussion

Across four studies (N = 1025), I 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 pseudo-profound bullshit receptivity was robust in that it was largely unaffected when controlling for potential cognitive and metacognitive mediators thought to underlie this association. Importantly,

I found evidence that people high in persuasive bullshitting appear unable to distinguish superficial profoundness (i.e., a statement simply "sounding profound") from inherent profoundness (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, I expand upon these findings and suggest some potentially fruitful directions for future research.

3.6.1 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.

I 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 I 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 Dunning-Kruger-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.

3.6.2 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 my 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.

3.6.3 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.

3.6.4 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. My primary aim was to examine the extent to which bullshitting frequency is associated with susceptibility to falling for bullshit. Overall, I 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

Baron, J. (2019). Actively open-minded thinking in politics. *Cognition, 188*, 8-18. https://doi.org/10.1016/j.cognition.2018.10.004

Bryant, F. B., & Yarnold, P. R. (1995). Principal-components analysis and exploratory and confirmatory factor analysis. In L. G. Grimm & P. R. Yarnold (Eds.), *Reading and understanding multivariate statistics* (pp. 99-136). Washington, DC, US: American Psychological Association.

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

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

Cassam, Q. (2018). Epistemic insouciance. *Journal of Philosophical Research, 43*, 1-20. DOI: 10.5840/jpr2018828131

Cheung, H., Siu, T. C., Chen, L. (2015). The roles of liar intention, lie content, and theory of mind in children's evaluation of lies. *Journal of Experimental Child Psychology, 132*, 1-13.

Chmielewski, M., & Kucker, S. C. (2019). An MTurk Crisis? Shifts in data quality and the impact on study results. *Social Psychological and Personality Science*. https://doi.org/10.1177/1948550619875149

Cillizza, C. (2019, October 1). *Mike Pompeo gave a very evasive answer when asked about the Ukraine call.* CNN Politics. 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

DePaulo, B. M., Kashy, D. A., Kirkendol, S. E., Wyer, M. M., & Epstein, J. A. (1996). Lying in everyday life. *Journal of Personality and Social Psychology, 70*(5), 979-995. doi:10.1037/0022-3514.70.5.979

De Vries, R. E., Lee, K., & Ashton, M. C. (2008). The Dutch HEXACO personality inventory: Psychometric properties, self-other agreement, and relations with psychopathy among low and high acquaintanceship dyads. *Journal of Personality Assessment, 90*(2), 142–151. DOI: 10.1080/00223890701845195

De Vries, R. E., Hilbig, B. E., Zettler, I., Dunlop, P. D., Holtrop, D., Lee, K., & Ashton, M. C. (2018). Honest people tend to use less—not more—profanity: Comment on Feldman et al.'s (2017) Study 1. *Social Psychological and Personality Science, 9*(5), 516–520. https://doi.org/10.1177/1948550617714586

De Vries, R. E., Zettler, I., & Hilbig, B. E. (2014). Rethinking trait conceptions of social desirability scales: Impression management as an expression of honesty-humility. *Assessment, 21*(3), 286–299. https://doi.org/10.1177/1073191113504619

Erlandsson A., Nilsson A., Tinghög G., Västfjäll D. (2018). Bullshit-sensitivity predicts prosocial behavior. *PLOS ONE 13*(7): e0201474. https://doi.org/10.1371/journal.pone.0201474

Evans, A. M., Sleegers, 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.

Feldman, R. S., Forrest, J. A., & Happ, B. R. (2002) Self-presentation and verbal deception: Do self-presenters lie more? *Basic and Applied Social Psychology, 24*(2), 163-170. DOI: 10.1207/S15324834BASP2402_8

Frankfurt, H. (2005 [1986]). *On bullshit*. Princeton, NJ: Princeton University.

Frankfurt, Harry G. 2006. *On truth*. New York: Alfred A. Knopf.

Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives, 19*(4), 25–42.

Giles, R. M., Rothermich, K., & Pell, M. D. (2019). Differences in the evaluation of prosocial lies: A cross-cultural study of Canadian, Chinese, and German adults. *Frontiers in Communication: Language Science*, https://doi.org/10.3389/fcomm.2019.00038

Gligorić, V., Feddes, A., & Doosje, B. (2020, October 27). Political bullshit receptivity and its correlates: a cross-cultural validation of the concept. *PsyArxiv*, https://psyarxiv.com/u9pe3/

Goldberg, L. R., Johnson, J. A., Eber, H. W., Hogan, R., Ashton, M. C., Cloninger, C. R., & Gough, H. C. (2006). The International Personality Item Pool and the future of public-domain personality measures. *Journal of Research in Personality, 40*, 84-96.

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.

Halevy, R., Shalvi, S., & Verschuere, B. (2014). Being honest about dishonesty: Correlating self-reports and actual lying. *Human Communication Research, 40*, 54– 72. https://doi.org/10.1111/hcre.12019

Handel, M. I. (1982). Intelligence and deception. *The Journal of Strategic Studies*, *5*(1), 122-154.

Harman, J. P., Hansen, C. E., Cochran, M. E., & Lindsey, C. R. (2005). Liar, Liar: Internet faking but not frequency of use affects social skills, self-esteem, social anxiety, and aggression. *CyberPsychology & Behavior*, *8*(1), 1-6. http://doi.org/10.1089/cpb.2005.8.1

Hart, C. M., Ritchie, T. D., Hepper, E. G., & Gebauer, J. E. (2015). The Balanced Inventory of Desirable Responding Short Form (BIDR-16). *SAGE Open*. https://doi.org/10.1177/2158244015621113

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

Hutcheson G., & Sofroniou N. (1999). *The multivariate social scientist: introductory statistics using generalized linear models.* London: Sage Publication.

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

Jones, D. N., & Paulhus, D. L. (2017). Duplicity among the dark triad: Three faces of deceit. *Journal of Personality and Social Psychology, 113*(2), 329-342. http://dx.doi.org/10.1037/pspp0000139

Judge, T. A., Erez, A., Bono, J. E., & Thoresen, C. J. (2003). The core self-evaluations scale: Development of a measure. *Personnel Psychology, 56*(2), 303-331. DOI 10.1111/j.1744-6570.2003.tb00152.x

Kaiser, H. F. (1974). An index of factorial simplicity. *Psychometrika*, *39*, 31–36.

Kajonius, P. J. (2014). Honesty–Humility in contemporary students: Manipulations of self-image by inflated IQ estimations. *Psychological Reports: Sociocultural Issues in Psychology*, *115*(1), 311-355.

Kim, H. (2013). Statistical notes for clinical researchers: Assessing normal distribution (2) using skewness and kurtosis. *Restorative Dentistry and Endodontics 38*, 52–54. https://doi.org/10.5395/rde.2013.38.1.52

Kimbrough, S. (2006). "On letting it slide." In G. L. Hardcastle & G. A. Reisch (Eds.), *Bullshit and philosophy: Guaranteed to get perfect results every time* (pp. 1-18). Chicago: Open Court.

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. doi: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*, 1 - 10.

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. (2021a). The Bullshitting Frequency Scale: Development and psychometric properties. *British Journal of Social Psychology, 60*(1), 248-270. https://doi.org/10.1111/bjso.12379

Littrell, S., Risko, E. F., & Fugelsang, J. A. (2021b). "You can't bullshit a bullshitter" (or can you?): Bullshitting frequency predicts receptivity to various types of misleading information. *British Journal of Social Psychology,* Advance online publication. https://doi.org/10.1111/bjso.12447

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.

Lupoli, M. J., Jampol, L., & Oveis, C. (2017). Lying because we care: Compassion increases prosocial lying. *Journal of Experimental Psychology: General, 146*(7), 1026-1042. http://dx.doi.org/10.1037/xge0000315

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

Martin, D. D., & Wilson, J. L. (2011). A Tool in the Kit: Uses of bullshitting among millennial workers. *The Qualitative Report, 16*(5), 1245-1269. Retrieved from https://nsuworks.nova.edu/tqr/vol16/iss5/3

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

Meibauer, J. (2018). The linguistics of lying. *Annual Review of Linguistics, 4,* 357-375. https://doi.org/10.1146/annurev-linguistics-011817-045634

Michels, M., Molz, G., & Maas genannt Bermpohl, 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

Nunally, J. C. (1978). *Psychometric theory* (2nd ed.). New York: McGraw-Hill.

Paulhus, D. L. (1991). Measurement and control of response bias. In J. P. Robinson, J. R. Shaver, & L. S. Wrightsman (Eds.), *Measure of personality and social psychological attitudes* (pp. 17-59). San Diego, CA: Academic Press.

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/zhjkp

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

Peterson, C., & Seligman, M. E. P. (2004). *Character strengths and virtues: A handbook and classification*. Oxford, United Kingdom: Oxford University Press.

Petrocelli, J. V. (2018). Antecedents of bullshitting. *Journal of Experimental Social Psychology, 76,* 249–258. https://doi.org/10.1016/j.jesp.2018.03.004

Price, E., Ottati, V., Wilson, C., & Kim, S. (2015). Open-minded cognition. *Personality and Social Psychology Bulletin, 41*(11), 1488-1504. DOI: 10.1177/0146167215600528

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.

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.

Roets, A. & Van Hiel, A. (2011). Item selection and validation of a brief, 15-item version of the need for closure scale. *Personality and Individual Differences, 50*, 90-94. https://doi.org/10.1016/j.paid.2010.09.004

Ruffle, B. J., & Tobol, Y. (2017). Clever enough to tell the truth. *Experimental Economics, 20*(1), 130-155. http://dx.doi.org/10.1007/s10683-016-9479-y

Serota, K. B., & Levine, T. R. (2015). A few prolific liars: Variation in the prevalence of lying. *Journal of Language and Social Psychology, 34*(2), 138–157. https://doi.org/10.1177/0261927X14528804

Spearman, C. (1904). The proof and measurement of association between two things. *The American Journal of Psychology, 15*(1), 72–101. DOI: 10.2307/1412159

Spicer, A. (2013). Shooting the shit: the role of bullshit in organisations. *M@n@gement, 16*(5), 653-666. doi:10.3917/mana.165.0653.

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

Stokke, A., & Fallis, D. (2017). Bullshitting, lying, and indifference toward truth. *Ergo, 4*(10), 277-309. http://dx.doi.org/10.3998/ergo.12405314.0004.010

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.

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

Williams B. (2002). *Truth and Truthfulness: An Essay in Genealogy*. Princeton, NJ: Princeton University Press.

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

Wright, E. J., White, K. M., & Obst, P. L. (2016). Facebook false self-presentation behaviors and negative mental health. *Cyberpsychology, Behavior, and Social Networking*, *21*(1), 1-10. http://doi.org/10.1089/cyber.2016.0647

Zettler, I., Hilbig, B. E., Moshagen, M., & de Vries, R. E. (2015). Dishonest responding or true virtue? A behavioral test of impression management. *Personality and Individual Differences*, *81*, 107-111. http://dx.doi.org/10.1016/j.paid.2014.10.007

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*

Zwick, W. R., & Velicer, W. F. (1982). Factors influencing four rules for determining the number of components to retain. *Multivariate Behavioral Research*, *17*(2), 253-269. http://dx.doi.org/10.1207/s15327906mbr1702_5

Appendix

A: Bullshitting measures

Bullshitting Frequency Scale (BSF)

Instructions: *On the following page, you will be given a number of statements that describe various situations people often encounter in their day-to-day lives when interacting and communicating with others.*

On a scale from "Never" to "All the time," please indicate how frequently you embellish, exaggerate, or otherwise stretch the truth when interacting with other people. You do not need to recall specific instances, though that might be helpful to you; just think about how often you do these things (or how likely you are to do them) in general.

It's important that we get accurate information about real human behaviour, so please respond honestly. Your responses are completely confidential.

1 = Never; 2 = Rarely; 3 = Occasionally/Sometimes; 4 = Frequently; 5 = A lot / All the time

In my daily life, I embellish, exaggerate, or otherwise stretch the truth:

1. When I want to impress the person or people I'm talking to.
2. When I want others to see me as more intelligent or knowledgeable.
3. When I want to contribute to a conversation or discussion even though I'm not well-informed on the topic.
4. By pretending to know more about a topic than I actually do.
5. When I'm trying to fit in better or be more accepted by the person or people I'm interacting with.
6. When I know it will be easy to get away with it.
7. When I want the thing(s) I'm talking about to sound more interesting or exciting.
8. When I'm trying to persuade someone to change their mind or agree with what I'm saying.
9. When being fully honest would be harmful or embarrassing to me or someone else.
10. When a direct answer might get me in trouble.
11. When I don't want to tell someone what I really think.
12. When a direct answer would hurt another person's feelings.

*Items should be presented in randomized order. Thus far, we have no data suggesting significant differences in response patterns from presenting them individually versus matrix format. Scoring information on page 2.

Scoring instructions

Persuasive bullshitting **subscale (BSFp)**: Items 1 thru 8. Calculate MEAN.
Evasive bullshitting **subscale (BSFe)**: Items 9 thru 12. Calculate MEAN.

Depending on the goals of the study, *for analyses using linear regression models, we recommend entering both subscales (BSFp and BSFe) as separate predictors* to account for overlapping variance (as individuals tend to engage in both types of BSing in their daily lives), rather than using an overall bullshitting (BSF) score as a single predictor.

Additionally, when examining associations with other variables, *we recommend calculating partial correlations controlling for each subscale separately* (see Littrell, Risko, & Fugelsang, 2020).

We suggest these methods of analysis to account for data suggesting that associations with each subscale for some variables (e.g., cognitive ability, cognitive reflection, bullshit receptivity, overclaiming, anti-social lying) are significantly different and/or trend in opposite directions, which can render an overall bullshitting score counterproductive.

Bullshitting social decision-making task

The following novel social decision-making task was designed for Study 5. Participants were presented with four vignettes and asked to rate the likelihood that they were respond with each of the three options. Response options were presented in randomized order.

1) Imagine that **you** are in the story. After you finish reading the story, answer the questions below.

   Your romantic partner suddenly gets a drastic haircut that he/she seems to really like. You do not like it at all and think that it is a really bad look, but you know that he/she is very sensitive about their appearance. Your partner asks you if you like it.

Rate how likely you would respond with the following if you were in this situation:

In this situation, would you say:	Definitely not	Probably not	Maybe	Probably, yes	Yes, definitely
"*I don't like it."* (truthful)	○	○	○	○	○
"*I like it!"* (lie)	○	○	○	○	○
"*Oh, wow that's a big change for you!"* (evasive)	○	○	○	○	○

*Parentheticals are there to identify the response category. These were not shown to participants.

2) Imagine that **you** are in the story. After you finish reading the story, answer the questions below.

You attend a friend's birthday party that they're hosting at their apartment. Though the food was good and you got to see old friends, you found the party to be pretty boring and you know that most other people who were there were also quite bored. As you grab your coat to leave, your friend who threw the party walks up and asks if you had fun.

Rate how likely you would respond with the following if you were in this situation:

In this situation, would you say:	Definitely not	Probably not	Maybe	Probably, yes	Yes, definitely
"*Not really."* (truthful)	○	○	○	○	○
"*Yes, it was fun!"* (lie)	○	○	○	○	○
"*Oh, the food was delicious! Did you make it?"* (evasive)	○	○	○	○	○

*Parentheticals are there to identify the response category. These were not shown to participants.

3) Imagine that **you** are in the story. After you finish reading the story, answer the questions below.

You receive a holiday gift from your favorite relative. After opening it, you realize that you really don't like it. Though you feel that they probably put a lot of thought into selecting the gift, it isn't at all something you would ever buy for yourself and you feel that you'll either throw it away or try to return it to the store. Your relative smiles and asks if you like it.

Rate how likely you would respond with the following if you were in this situation:

In this situation, would you say:	Definitely not	Probably not	Maybe	Probably, yes	Yes, definitely
"*Not really."* (truthful)	○	○	○	○	○
"Yes! Thank you! (lie)	○	○	○	○	○
"*Oh, wow! This was really thoughtful of you."* (evasive)	○	○	○	○	○

*Parentheticals are there to identify the response category. These were not shown to participants.

4) Imagine that **you** are in the story. After you finish reading the story, answer the questions below.

   After seeing a job posting at work, you decide to apply for a new position within the company in a different department. Although you love the field that you're in and you always get your projects finished on time, you generally hate your job, mostly because you really dislike your current boss. Although your boss is very organized, he/she is overly demanding, treats the employees poorly, and is just an all-around jerk.

   Your application is accepted and you go in for an interview. During the interview, the hiring manager asks you how you like working for your current supervisor in the other department.

Rate how likely you would respond with the following if you were in this situation:

In this situation, would you say:	Definitely not	Probably not	Maybe	Probably, yes	Yes, definitely
"*I'm not a fan."* (truthful)	○	○	○	○	○
"I like him/her. It's been a good experience." (lie)	○	○	○	○	○
"*He/she is very organized and ensures that everyone stays on task."* (evasive)	○	○	○	○	○

*Parentheticals are there to identify the response category. These were not shown to participants.