

The Bullshitting Frequency Scale: Development and psychometric properties

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

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

CONTACT: Shane Littrell, sslittrell@uwaterloo.ca

**This is an original manuscript / preprint of an article published by Wiley in
The British Journal of Social Psychology, on April, 18, 2020, available online at:**

<https://onlinelibrary.wiley.com/doi/full/10.1111/bjso.12379>

When citing, please list and cite the published version.

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, we present results from four studies (focusing on 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. 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. These results represent an important step forward 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 everyday bullshitting.

Key words: bullshitting, bullshit, lying, persuasion, overclaiming, analytic thinking

Introduction

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).

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. Therefore, 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 off of 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 echoes 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 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 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).

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).

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 algorithm¹. 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, Secarǎ, 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

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.

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 *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.

Present investigation

Here, we report four studies focused 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, we 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.

Study 1

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), we 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, we then examined bivariate and partial associations of the scale with various individual difference measures of related constructs, the results of which we present in Study 1b.

Study 1a – Scale creation

Method

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). 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 [removed for blinded review]. 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 us with data for 361

participants in the final analysis (222 male, 137 female, 2 intersex or prefer not to answer, $M_{\text{age}} = 36.40$, $SD_{\text{age}} = 11.26$, 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.

Procedure

After indicating consent and answering demographic questions (i.e., age, biological sex, and level of education), we presented participants 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.

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, $z(\text{skew}) = 6.77$. 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. 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).² To reduce scale size and ensure high factor reliability, redundant items and those with factor loadings < .500 were eliminated. A second principal components analysis was conducted which yielded the same pattern of factor loadings, therefore we retained this 12-item iteration as the final version of the scale (Table 1). Macdonald's scale reliabilities were strong for the full, 12-item scale ($\omega = .93$).

Confirmatory analysis

We 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). Factor loading plots and additional fit indices for both models can be found in the supplementary materials.

² Pattern matrix loadings for all items after rotation are listed in Table S1 of the supplementary materials.

Factor labelling

The items clustering around Factor 1 suggest that it represents 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), Petrocelli (2018), and Stokke and Fallis (2017), and we have labelled this factor *persuasive bullshitting*. Cronbach's alpha scale reliability for the *persuasive* factor was strong ($\alpha = .92$).

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) and elements of Mears (2002) views of bullshitting, we have labelled this factor *evasive bullshitting*. Cronbach's alpha scale reliability for this factor was also strong ($\alpha = .81$).

(insert Table 1)

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).

Study 1b – Associations with related constructs

We next examined correlations between our *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, we would also want it 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.

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-related measures

Participants completed the *Integrity/Honesty/Authenticity* scale from the IPIP version of the

Values in Action scale and the IPIP's version of the HEXACO *sincerity* scale (Goldberg et al., 2006; Peterson & Seligman, 2004). 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). Social desirability motivations were measured using the *self-deceptive enhancement* and *impression management* subscales of the Balanced Inventory of Desirable Responding (BIDR; Paulhus, 1991).

Cognitive ability

Numeracy and verbal intelligence were measured using a 10-item version of the General Risk and Numeracy Scale (Lipkus, Samsa, & Rimer, 2001) and a 10-item version of the "Wordsum" vocabulary test (Thorndike, 1942; Malhotra, Krosnick, & Haertel, 2007). Scale scores from both were combined to calculate a mean *cognitive ability* score.

Results

Table 2 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 (e.g., Müller & Moshagen, 2019). Therefore, while we have reported these associations in Table 2, they will not be discussed further (see supplementary materials for a full discussion).

(insert Table 2)

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, we 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, $r(358) = -.11, p = .04$.

Discussion

Consistent with our 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).

Study 2 – Bullshitting vs Lying

To further validate the BSF, in Study 2 we sought to discriminate the measurement of “everyday bullshitting” from that of “everyday lying” while also confirming its factor structure. Therefore, we 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*, we 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 1 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; Wright, Berry, Bird, 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, we 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.

Method

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). Our 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, $M_{\text{age}} = 37.74$, $SD_{\text{age}} = 11.72$, Bachelor's degree or higher = 54.8%).

Procedure

We preregistered our 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.

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).

Results

We first conducted a series of confirmatory factor analyses to both confirm the BSF’s factor structure as well as ensure that the BFS and LiES scales were distinguishable at the factor level. We next calculated descriptive statistics as well as bivariate and partial Pearson’s *r*-values for all study variables (Table 3) and compared correlation coefficients using Fishers *r*-to-*z* transformations. We first report the factor analyses, followed by the bivariate results.

Confirmatory factor analyses

A confirmatory factor analysis (CFA) using JASP (v0.11.1.0) 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). We 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 (see supplementary materials for fit indices and path diagrams for each model).

(insert Table 3)

Bivariate correlations

Scores on the overall BSF and LiES scales (Table 3) 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, $|z| = 2.26, p = .01$.

BSF, $r(315) = -.22, p < .001$, and LiES, $r(315) = -.37, p < .001$, scores 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, $|z| = 0.38, p = .35$.

Partial correlations

We 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, $|z| = 4.43, p < .01$.

Discussion

Results from Study 2 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 we 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.

Study 3 – Persuasive and evasive bullshitting

Our goals in Study 3 were twofold. First, we wanted to test four reworded BSF items that we felt might enhance the content validity and factor structure of each subscale. Additionally, we sought to further distinguish *persuasive* and *evasive bullshitting*. Our 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, we presented participants in Study 3 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. We found some evidence in Study 1 that persuasive bullshitting was positively related to an overclaiming task while evasive was not, so we 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 actually 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.

Our 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.

Participants and procedure

Preregistration for Study 3 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.

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. We 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 supplementary materials). 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.

Results

Per our 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, $M_{\text{age}} = 36.28$, $SD_{\text{age}} = 10.89$, Bachelor's degree or higher = 57.3%). Bivariate and partial correlations, as well as descriptive statistics, for all study variables are listed in Table 5.

Confirmatory analysis 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 (see Table 4).

We 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). Factor loading plots and additional fit indices for both models can be found in the supplementary materials.

(insert Table 4)

Bivariate correlations

As our hypotheses were concerned specifically with distinguishing the BSF subscales, our discussion of the results will focus on the subscales only, though data for the overall BSF scale is listed in Table 5. *Persuasive* and *evasive bullshitting* were correlated at the bivariate

level $r(298) = .62, p < .001$. Additionally, *persuasive bullshitting* was significantly and positively related to overclaiming $r(298) = .18, p < .001$, proposition overclaiming, $r(298) = .15, p = .01$, and lie responses on the social decision task, $r(298) = .40, p < .001$, and significantly and negatively related to cognitive ability, $r(298) = -.31, p < .001$. *Evasive bullshitting* was significantly and positively related to lie responses, $r(298) = .39, p < .001$, and evasive responses, $r(298) = .12, p = .04$, on the social decision task and significantly and negatively related to cognitive ability, $r(298) = -.16, p < .001$.

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, $r(297) = -.17, p < .001$.

(insert Table 5)

Discussion

Consistent with our predictions, the results from Study 3 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 2), 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 2, 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.).

General Discussion

In the present investigation, we 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).

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 we 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, we 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.

Being honest about bullshitting

It is also important to consider the extent to which we can be confident that participants responded honestly to items on the BSF and were not, in fact, bullshitting the researchers. Here, we 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, we 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 our 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 our 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 we used in our study, we feel confident that our data represent overall honest and accurate responding patterns from participants.

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).

References

- Carson, T. (2016). Frankfurt and Cohen on bullshit, bullshitting, deception, lying, and concern with the truth of what one says. *Pragmatics & Cognition* 23(1), 54-68. <https://doi.org/10.1075/pc.23.1.03car>
- Cassam, Q. (2018). Epistemic insouciance. *Journal of Philosophical Research*, 43, 1-20. <https://doi.org/10.5840/jpr2018828131>
- Čavoјová, V., Secarǎ, 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>
- 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. <https://doi.org/10.1016/j.jecp.2014.12.002>
- Chmielewski, M., & Kucker, S. C. (2019). An MTurk Crisis? Shifts in data quality and the impact on study results. *Social Psychological and Personality Science* 11(4), 464– 473. <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>

- 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., 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. <https://doi.org/10.1080/00223890701845195>
- 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>
- 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
- Dunlop, P. D., Bourdage, J. S., de Vries, R. E., McNeill, I. M., Jorritsma, K., Orchard, M., ... Choe, W.-K. (2019). Liar! Liar! (when stakes are higher): Understanding how the overclaiming technique can be used to measure faking in personnel selection. *Journal of Applied Psychology*. Advance online publication. <https://doi.org/10.1037/apl0000463>
- 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>

- 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. (1986). On bullshit. *Raritan Quarterly Review, 6*, 81-100.
- 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, 4*(38), 1– 11.
<https://doi.org/10.3389/fcomm.2019.00038>
- 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*(1), 84-96. <https://doi.org/10.1016/j.jrp.2005.08.007>
- 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>
- 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. L., Jones, J. M., Terrizzi, J. A., & Curtis, D. A. (2019). Development of the Lying in Everyday Situations scale. *The American Journal of Psychology, 132*(3), 343– 352.

- 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>
- 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.
- 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.

- Litman, L., Robinson, J., & Abberbock, T. (2016). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 1 - 10.
- 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>
- Mears, D. P. (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. (2018). The linguistics of lying. *Annual Review of Linguistics*, 4, 357-375. <https://doi.org/10.1146/annurev-linguistics-011817-045634>
- Müller, S., & Moshagen, M. (2019). True virtue, self-presentation, or both?: A behavioral test of impression management and overclaiming. *Psychological Assessment*, 31(2), 181– 191. <https://doi.org/10.1037/pas0000657>
- Nunally, J. C. (1978). *Psychometric theory* (2nd ed.). New York: McGraw-Hill.
- Oliveira, C. M., & Levine, T. R. (2008). Lie acceptability: A construct and measure. *Communication Research Reports*, 25(4), 282– 288. <https://doi.org/10.1080/08824090802440170>

- 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.
- Paulhus, D. L., Harms, P. D., Bruce, M. N., & Lysy, D. C. (2003). The over-claiming technique: Measuring self-enhancement independent of ability. *Journal of Personality and Social Psychology*, *84*(4), 890– 904.
<https://doi.org/10.1037/0022-3514.84.4.890>
- 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. doi.org/10.3389/fpsyg.2013.00279
- Pennycook, G., & Rand, D. G. (2019). Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *Journal of Personality*. Advance online publication. <https://doi.org/10.1111/jopy.12476>
- 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
- 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.

- 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>
- Serota, K. B., Levine, T. R., & Boster, F. J. (2010). The prevalence of lying in America: Three studies of self-reported lies. *Human Communication Research*, 36(1), 2–25. <https://doi.org/10.1111/j.1468-2958.2009.01366.x>
- 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.
- 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>
- Thorndike, R. L. (1942). Two screening tests of verbal intelligence. *Journal of Applied Psychology*, 26(2), 128–135.
- 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.
- Williams B. (2002). *Truth and Truthfulness: An Essay in Genealogy*. Princeton, NJ: Princeton University Press.
- 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>

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>

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>

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

Tables

Table 1

Table 1
Pattern matrix factor loadings after rotation for each scale item with means and standard deviations

		1	2	<i>M</i>	<i>SD</i>
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, $r = .76$; * = replaced in Study 4.

Table 2

Table 2

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

	<i>M</i>	<i>SD</i>	α	<i>Bivariate</i>			<i>Partial</i>	
				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. ** $p < .01$; * $p < .05$

Table 3

Table 3

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

	<i>M</i>	<i>SD</i>	α	<i>Bivariate</i>						<i>Partial</i>		
				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; BSFe^b = Evasive bullshitting, controlling for persuasive.

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

Table 4

Table 4

Pattern matrix factor loadings after rotation for each scale item

	1	2	<i>M</i>	<i>SD</i>
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, $r = .59$

Table 5

Table 5

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

	<i>M</i>	<i>SD</i>	α	<i>Bivariate</i>			<i>Partial</i>	
				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. ** $p < .01$; * $p < .05$