Bullshit blind spots: The roles of miscalibration and information processing in bullshit detection

Shane Littrell¹ and Jonathan A. Fugelsang

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

This is an original manuscript / preprint of an article published by Taylor and Francis in *Thinking and Reasoning*, available online at:

<https://doi.org/10.1080/13546783.2023.2189163>

When citing, please list and cite the published version:

Littrell, S., & Fugelsang, J. A. (2023). Bullshit blind spots: The roles of miscalibration and information processing in bullshit detection. *Thinking and Reasoning*, Advance online publication. <https://doi.org/10.1080/13546783.2023.2189163>

¹ Correspondence concerning this article should be addressed to Shane Littrell, Department of Psychology, University of Waterloo, Waterloo, ON, Canada N2L 3G1, sslittrell@uwaterloo.ca

Abstract

The growing prevalence of misleading information (i.e., bullshit) in society carries with it an increased need to understand the processes underlying many people's susceptibility to falling for it. Here we report two studies ($N = 412$) examining the associations between one's ability to detect pseudo-profound bullshit, confidence in one's bullshit detection abilities, and the metacognitive experience of evaluating potentially misleading information. We find that people with the lowest (highest) bullshit detection performance overestimate (underestimate) their detection abilities and overplace (underplace) those abilities when compared to others. Additionally, people reported using both intuitive and reflective thinking processes when evaluating misleading information. Taken together, these results show that both highly bullshit-receptive and highly bullshit-resistant people are largely unaware of the extent to which they can detect bullshit and that traditional miserly processing explanations of receptivity to misleading information may be insufficient to fully account for these effects.

Key words: bullshit receptivity, misinformation, overconfidence, dual process theory, bias blind spot

Introduction

A recent global survey on the spread of online misinformation conducted by international marketing research firm *Ipsos* found that nearly 86% of all respondents reported that they have been exposed to misinformation on social media with roughly 90% of those admitting to initially being duped by at least one misleading story (Simpson, 2019). Moreover, some researchers have mapped how information is shared online, finding that false and misleading information spreads faster and to a significantly larger audience than the truth across all categories of information that were studied (Juul & Ugander, 2021; Vosoughi et al., 2018).

Given this current misinformation milieu, a growing body of work has emerged examining factors associated with the transmission and reception of misleading information, with particular attention paid to so-called “fake news” and various types of scientific, political, and pseudo-profound “bullshit” (e.g., Evans et al., 2020; Littrell, et al, 2021b; Pennycook & Rand, 2020; Petrocelli, 2021). Much of this research has focused on identifying specific cognitive factors common among people more susceptible to misleading information, finding negative relations between credulity for various types of bullshit with analytic thinking, open-minded cognition, intelligence, and metacognitive insight, to name a few (Bago, et al., 2020; De keersmaecker & Roets, 2017; Littrell et al., 2021b; Pennycook & Rand, 2020). However, two potentially critical factors in regulating one’s susceptibility to misleading information, and those we seek to shed light on here, are the extent to which a person is aware of their own receptivity to bullshit, and their perceptions of the thinking processes they use to detect and resist its persuasive influence.

Receptivity to bullshit

We define bullshit here as information constructed with a carefree indifference for conveying truth, accuracy, clarity, or meaning that is often used to impress, persuade, or otherwise mislead others (Cohen, 2012; Frankfurt, 2006; Littrell et al., 2021b; Pennycook et al., 2015a; Reisch, 2006). The propensity to give inflated positive judgments of qualities related to the persuasiveness (e.g., truthfulness, meaningfulness, likeability, etc.) of bullshit is known as one's *bullshit receptivity* (Evans et al., 2020; Littrell et al., 2021b; Pennycook et al., 2015a). Pennycook and Rand have speculated that bullshit receptivity may be driven by “reflexive open-mindedness,” which describes both a willingness to accept and a failure to detect epistemically-suspect claims (Pennycook & Rand, 2020).

Research into individual differences in bullshit receptivity was jump-started by the introduction of the Bullshit Receptivity Scale (BSR) which measures one's receptivity to pseudo-profound bullshit (Pennycook et al., 2015a). This measure asks participants to evaluate the profoundness of a series of statements. Half of the statements are well-known sayings and motivational quotes that were intended to convey a profound message or idea. The remaining statements were randomly constructed from a list of New Age buzzwords and jargon by a computer algorithm to be syntactically correct but semantically meaningless, resulting in statements constructed without an intention to convey deep meaning or truth (i.e., “pseudo-profound bullshit”). Similar measures have been devised for assessing individual differences in receptivity to the perceived truthfulness of scientific bullshit (Evans et al., 2020), the meaningfulness of financial bullshit (Kienzler et al., 2022), and the likeability of general bullshit (Čavojová et al., 2020).

Higher scores on the BSR have been found to reliably predict greater acceptance of a broad range of epistemically-suspect beliefs associated with decreased Type II engagement, such

as conspiracy theories, fake news, and scientific misinformation (Evans et al., 2020; Pennycook et al., 2015a, Pennycook & Rand, 2020; van Prooijen et al., 2022). Bullshit receptivity is also positively associated with sharing misinformation online and greater engagement in persuasive bullshitting (Čavojová et al., 2019; Evans et al., 2020; Littrell et al., 2021b).

Past research has explained higher bullshit receptivity via a dual process account (e.g., Pennycook et al., 2015a). That is, being more likely to fall for bullshit is thought to result from a greater reliance on intuitive Type I thinking processes coupled with decreased engagement of more reflective, Type II processes when evaluating information. Indeed, several studies have found that higher scores on the BSR are negatively associated with measures of analytic thinking, open-minded cognition, and intelligence (Čavojová et al., 2019; Littrell et al., 2021b; Pennycook et al., 2015a). Given these findings, bullshit receptivity scores are seen as a good index of susceptibility to a broad range of misinformation and decreased engagement in reflective thinking. However, the extent to which bullshit receptivity may be related to important metacognitive factors which might serve to override reflective engagement, such as overconfidence and metacognitive monitoring, remain unclear.

Overconfidence and reflective engagement

Overconfidence has been demonstrated in a variety of areas such as driving ability, general knowledge, and analytic reasoning skills, where those who possess the lowest levels of competence in specific domains often appear blissfully unaware of their lack of prowess in those areas (Kruger & Dunning, 1999; Macenczaka et al., 2016; Roy & Liersch, 2013). However, only recently have researchers begun to consider the large-scale repercussions that overconfidence may have on the ways people process information, including how it might contribute to susceptibility to misinformation (e.g. Salovich & Rapp, 2020; Sheldon et al., 2014). For

example, recent work has shown that overconfidence in one's cognitive and decision-making abilities may function to override engagement in certain analytic thinking processes in favour of incorrect, intuitive responses (Littrell et al., 2020). Additionally, some research has found that those who score lowest on measures of analytic thinking are not only overconfident in their reasoning skills but ironically tend to self-report higher levels of dispositionally analytic thinking, despite their low objective performance on such tasks (Pennycook et al., 2017). This suggests that overconfident reasoners may have diminished metacognitive insight for their own engagement in specific thinking processes. Moreover, a growing body of work has suggested that people who are more receptive to bullshit are less likely to engage in analytic thinking (Evans et al, 2020; Pennycook et al., 2015a; Pennycook & Rand, 2020). If highly bullshit receptive people are prone to overestimating their actual bullshit detection ability, this may reflect a lack of insight into the thinking processes they utilize to evaluate and detect bullshit. Based on these findings, it stands to reason that many people who are especially receptive to bullshit may experience levels of confidence that are miscalibrated to their actual ability to detect it, overriding their engagement in reflective thinking when evaluating misleading information, thus rendering them largely unaware of their own susceptibility.

Bias blind spots

In addition to believing that their knowledge or performance on certain tasks is dramatically better than it objectively is, people who are overconfident also tend to believe that their knowledge or performance is measurably better than that of the average person on those same tasks; a phenomenon known as the better-than-average bias (BTA) or *overplacement* (Merkel & Weber, 2011; Moore & Schatz, 2017; Zell et al., 2020). Furthermore, BTA beliefs can extend to a person's views of their own ability to think critically and resist biases, including

the insistence that their self-assessments of these abilities are more accurate and objective than those of most other people (Scopelliti et al., 2015; Pronin et al., 2002). These types of BTA beliefs represent a self-perception error that has been dubbed the “bias blind spot,” (Pronin et al., 2002). Better-than-average effects tend to be statistically robust with medium-to-large effect sizes and have been found across a wide range of domains (Zell et al., 2020). Some researchers have suggested that self-enhancement underlies BTA beliefs, thus a bias blind spot could result from a desire to maintain positive self-regard (Pronin et al., 2002; Zell et al., 2020). Critically, a recent meta-analysis by Zell and colleagues (2020) reports a medium-sized association between BTA effects and self-esteem ($r = .34$) and suggests that self-enhancement is a major pillar of BTA and blind spot beliefs. Other researchers have theorized that bias blind spots are somewhat driven by naïve realism, the belief that one’s perceptions represent the true state of the world (Scopelliti et al., 2015). Additionally, blind spot biases may not be sufficiently attenuated by higher levels of cognitive sophistication (West et al., 2012). Thus, even those who are highly intelligent and analytically-minded are not immune to bias blind spots.

Although we are unaware of any previous work directly examining the associations between BTA beliefs and bullshit receptivity, recent research has found that bullshit receptivity is positively associated with overconfidence in one’s intellectual abilities (in general) as well as overconfidence in one’s performance on intelligence tasks (Littrell et al., 2021b). Additionally, bullshit receptivity is positively predicted by greater self-reported frequency of persuasive bullshitting, which has been found to be negatively related to self-regard and is itself often construed as an attempt at self-enhancement (Littrell et al., 2021a, 2021b). Thus, it could be the case that higher receptivity to misleading information is at least partially driven by overconfidence in one’s ability to detect it, and that this overconfidence is associated with

specific thinking processes and self-enhancing social comparisons. Together, such overconfidence could render highly bullshit receptive people not only less likely to successfully engage in the type of skeptical analytic thinking processes that would help them better identify and detect bullshit, but also largely oblivious to the fact that they are often casualties of its allure (Littrell et al., 2021b).

Present investigation

Based upon the prevailing research, we suggest that highly bullshit receptive individuals may experience: 1) a misalignment of confidence in their ability to detect bullshit relative to actual performance; 2) a “blind spot” biasing them to believe that they are significantly superior to others at detecting bullshit, and; 3) a lack of insight related to perceptions of their own thinking processes when evaluating misleading information. Here we report two studies ($N = 412$) examining the associations between bullshit detection, overconfidence, and the perceived thinking processes used in detection. Specifically, we look at general patterns underlying the putative association between bullshit detection confidence and actual bullshit detection ability. We also examine how these associations might be related to an individual’s beliefs about others’ bullshit detection abilities as well as the metacognitive and analytic strategies individuals perceive that they use when detecting and evaluating bullshit. Data files for all studies are available here: <https://osf.io/ynufw/>

Study 1

In Study 1, we examine the extent to which perceived ability to detect bullshit (i.e., confidence) is associated with actual bullshit detection ability. Consistent with previous research (Kruger & Dunning, 1999), we expect that, in general, people who demonstrate low bullshit

detection ability will be overconfident in their ability relative to those with higher bullshit detection ability. We also assess people's beliefs about their own bullshit detection ability relative to their beliefs about the performance of others. That is, we seek to determine whether people who are more receptive to pseudo-profound bullshit experience a "bias blind spot" in regard to perceptions of their own bullshit detection abilities compared to their perceptions of others' abilities.

Method

Participants

We recruited 223 adult participants from the United States and Canada from Amazon's Mechanical Turk (MTurk) participant pool using the crowdsourcing platform, CloudResearch (Litman, Robinson, & Abberbock, 2016). Given the recent increased focus on issues related to the quality and reliability of online data, we implemented several measures to ensure high data quality (e.g., Chmielewski et al., 2020; Ophir et al., 2020). Firstly, we restricted recruitment to participants who had completed a minimum of 500 surveys and had an MTurk HIT approval rating of at least 99%. Importantly, we also utilized two pre-survey eligibility questions designed to help decrease the number of fake responders gaining access to our survey (e.g., survey farmers using algorithmic "bots" to generate survey responses). These "botchas" ("bot gotchas") successfully prevented 23 potential low-quality responders from accessing our survey (a discussion of the botchas can be found in the Supplementary Materials).

Of the 223 who passed our eligibility criteria, data for 12 participants were removed for failing to pass reCAPTCHA v3 bot detection protocols or for failing an attention check, leaving us with data for 211 participants for analysis (116 males, 90 females, 5 prefer not to answer, M_{age}

= 38.87, $SD_{\text{age}} = 11.92$, Bachelor's degree or higher = 61%). An a priori power analysis was conducted with the goal of achieving at least .80 power to detect a medium effect (Cohen's $f = .25$) at $\alpha = .05$, indicating that we would need a sample of 180, which our sample exceeded (g*power; Faul, Erdfelder, Buchner, & Lang, 2009). Participants were paid \$2.00 USD for the roughly 15-minute study.

Procedure

The survey was hosted using Qualtrics Online Survey Platform (<https://www.qualtrics.com>). Those who passed both pre-screen eligibility questions (botchas) read an informed consent form and, after agreeing to participate, answered three demographic questions (i.e., age, biological sex, and level of education). Participants then completed the remainder of the survey which included the following measures noted in the materials section below. Participants also completed measures of fluid intelligence and bullshitting frequency, presented before the bullshit detection task and in random order. However, these measures were included for reasons peripheral to the focus of the current paper, so they will not be discussed here.

Materials

Bullshit detection

Participants completed the Bullshit Receptivity Scale (BSR; Pennycook et al., 2015a) in which they evaluated the profoundness of 20 statements. The BSR was originally designed using a 1 (not at all profound) to 5 (Very profound) scale, which we modified into a forced-choice scale of “Profound” and “Not profound” according to the following instructions (bolded passages were also bolded for participants):

We are interested in how people experience the profound. Statements that are profound communicate messages or ideas “of deep meaning; of great and broadly inclusive significance.”

The following items are a series of statements taken from relevant websites. Several of the statements were made with the intention to sincerely communicate a profound message or idea while others were not and may only seem profound.

We would like you to identify which statements were made to convey a profound message or idea and which were not using these criteria:

***"Not profound"** for statements that convey a message or idea that is not deep or profound.*

***"Profound"** for statements that convey a message or idea that is at least somewhat to definitely deep or profound.*

We used this forced-choice rating scheme (rather than the usual 5-point rating scale) to decrease the potential for ambiguity in profoundness evaluations. Ten of these statements were randomly generated² – yet grammatically correct – sentences constructed from abstract pseudo-profound buzzwords (e.g., “Perceptual reality transcends subtle truth”). That is, they were statements that were not, per the instructions, “designed with the intention to sincerely communicate a profound message or idea.” The remaining 10 items were a collection of intentionally profound/motivational quotes from influential historical figures taken from a number of sources (e.g., “Your teacher can open the door, but you must enter by yourself.”). All items were presented in a randomized order.

From this, we calculated a *bullshit detection* score (percent-accuracy) based on the number of times profound items had been correctly identified as “profound” and non-profound items had been correctly identified as “not profound,” according to the following formula:

² To eliminate redundant phrasings and sentence structures, we replaced 5 BSR items with new ones created using the same online generator (<https://sebpearce.com/bullshit/>) that Pennycook and colleagues (2015) used when constructing the original BSR items (see Supplementary Materials).

Total correct = (# of correct motivational quote hits) + (# of correct pseudo-profound rejections)

$$\text{Bullshit detection accuracy} = \left(\frac{\text{Total correct}}{20} \right) \times 100$$

We used this conversion so that the scale of our bullshit detection variable would match that of participants' confidence estimates, allowing for more straightforward calculations of the overconfidence variables. Importantly, to ensure that our "bullshit detection accuracy" score captures *bullshit sensitivity*, also sometimes referred to as *bullshit discernment* (i.e., one's ability to distinguish bullshit from non-bullshit), we also calculated a sensitivity (discernment) score for each participant by subtracting the proportion of "false alarms" (i.e., incorrectly rating pseudo-profound bullshit statements as "profound") from "hits" (i.e., correctly rating intentionally profound/motivational statements as "profound"). We then confirmed the equivalence of the *bullshit detection accuracy* and *bullshit sensitivity* variables via their bivariate correlation, $r = 1.0$.

Confidence ratings

Data for two types of estimated detection confidence were collected: *estimates of own performance* and *estimates of others' performance*. To assess confidence in one's own performance on the bullshit detection task, each participant estimated their performance immediately afterwards by rating the following item using a 0 to 100 sliding scale. The purpose of the task (i.e., to distinguish statements that were created to be intentionally profound from those that were not created to be intentionally profound) was again emphasized in the instructions for clarity:

“In the previous section, you were asked to classify 20 items. Some statements were intended to be profound, and others were not intended to be profound. What % of the statements do you think you correctly classified? That is, you correctly selected "Profound" for the intentionally profound items and "Not Profound" for the items that were not intentionally profound.”

Individual *estimated-self* ratings were then used to calculate an *overestimation* index for each participant using the following equation from Moore and Healy (2008):

$$\text{Overestimation} = \text{Individual's estimated score} - \text{Individual's actual score}$$

Participants also used a 0 to 100 sliding scale to estimate what they thought the average performance (in % correct responses) was of all other participants on the bullshit detection task (“*What is the average % of the statements that you think other participants correctly identify?*”). Ratings for these items were used to calculate *overplacement*, a self-serving bias which reflects the belief that one is better than the average participant, corrected for the degree to which one actually *is* better (or worse) than the average participant. As such, it is not an absolute measure of how much better/worse a person believes they performed than the average person, but a bias score based on the degree to which such a belief is warranted. Scores greater than 0 reflect *overplacement* and those less than 0 reflect *underplacement*. Overplacement scores for each participant were calculated according to the following equation from Moore and Healy (2008):

$$\begin{aligned} \text{Overplacement} = & (\text{Estimate of own score} - \text{Estimate of other's score}) \\ & - (\text{Actual own score} - \text{average accuracy of all scores}) \end{aligned}$$

Self-reported confidence in overall bullshit detection ability

Participants also self-reported their perceived confidence in their ability to detect bullshit in general. For this, participants rated the item, “*In comparison to other people, how would you rate your ability to detect bullshit (e.g., instances when someone is trying to bullshit you, misinformation online, etc.)?*” using a 0 (“Much worse than others”) to 100 (“Much better than others”) sliding scale.

Results and discussion

Based on procedures from Kruger and Dunning (1999), we split the sample into four groups based on quartiles for *bullshit detection accuracy* to test whether those who demonstrated lower bullshit detection were more overconfident than those who exhibited better bullshit detection. Mean scores for all variables of interest are listed in Table 1.

Table 1
Mean scores (standard deviations) by accuracy quartile

	<i>Quartile 1 (lowest)</i>	<i>Quartile 2</i>	<i>Quartile 3</i>	<i>Quartile 4 (highest)</i>
1 Actual accuracy	38.02 (8.23)	54.35 (4.10)	70.09 (4.10)	86.15 (5.67)
2 Estimated accuracy	66.13 (17.43)	63.11 (16.43)	57.24 (19.61)	65.69 (18.69)
3 Self-reported general BS detection ability	72.81 (18.34)	67.55 (17.41)	64.60 (20.07)	65.54 (17.16)
4 Overestimation	28.10 (18.99)	8.76 (16.99)	-12.85 (20.05)	-20.46 (18.26)
5 Overplacement	26.95 (15.21)	7.48 (14.00)	-11.20 (15.75)	-22.49 (11.80)

Note: $N = 211$. Quartiles listed from lowest to highest performing. *SDs* are in parentheses.

Overestimation and general BS detection overconfidence

We first sought to examine general patterns underlying the association between bullshit detection confidence and actual bullshit detection performance. Specifically, we were interested in examining whether the patterns revealed notable differences in levels of overconfidence or underconfidence for those individuals who are especially poor at detecting bullshit and those who are especially good at detecting it. To examine these possible differences, we first plotted mean estimated bullshit detection scores and mean accuracy scores by quartile (see Figure 1).

Figure 1. Miscalibration for estimated BS detection scores relative to actual BS detection accuracy (by quartile) in Study 1. The y-axis represents the actual/estimated percent correct for the 20 bullshit detection items. Error bars represent 95% confidence intervals.

A distinct pattern emerged such that those who performed the worst overestimated their performance while those whose scores were highest underestimated their performance. Crucially, this pattern somewhat differs from what is more commonly found in similar overconfidence studies. Whereas the typical pattern tends to reflect a greater confidence-accuracy gap for the

lowest quartile performers compared to the highest quartile, the present results show that participants in the bottom and top quartiles were approximately equally miscalibrated, though in opposite directions. Indeed, the extent to which the lowest performers were overconfident was similar to the degree of underconfidence for the highest performers. Moreover, participants across all accuracy quartiles tended to estimate their bullshit detection scores at roughly the same level, suggesting that both low and high performers are similarly unaware (though in opposite directions) of the extent of their ability to detect bullshit.

To explore this pattern further, we next examined differences in overestimation across the four performance groups. As our overestimation variable is a difference score between actual accuracy and estimated accuracy, we conducted a one-way between-groups ANOVA to compare differences in bullshit detection overestimation across quartile levels of actual bullshit detection accuracy (see Figure 2A), finding a statistically significant difference in means, $F(3,207) = 69.12, p < .001, \eta^2 = 0.50$. A post-hoc Tukey HSD test comparing overestimation among the four groups showed significant mean differences between three out of the four groups ($p < .001$). Specifically, those who scored lowest in bullshit detection (i.e., the bottom quartile, $M = 38.02, SD = 8.23$) overestimated their performance by an average of 28.10%. Those in the next lowest accuracy quartile ($M = 54.35, SD = 4.10$) fared better, overestimating their performance by an average of only 8.76%. Conversely, individuals in the higher accuracy quartiles tended to be underconfident in their bullshit detection performance. For example, those in the third quartile ($M = 70.09, SD = 4.10$) *underestimated* their bullshit detection score by an average of 12.85% while those in the top quartile ($M = 86.15, SD = 5.67$) underestimated their performance by an average of 20.46%. The mean difference in overestimation between the 3rd and 4th quartiles ($M = 7.61$), however, was not significant ($p = .17$).

Given these results, as well as a visual examination of Figure 1, the lowest performers appeared to be positively miscalibrated (i.e., overconfident) to a somewhat similar degree that the highest performers were negatively miscalibrated (i.e., underconfident). Therefore, we next calculated a *calibration index* for all participants by squaring their overestimation scores (i.e., bias²) to eliminate negative values (Schraw, 2009). This variable provides us with an index representing the absolute degree to which a person's confidence is misaligned with their accuracy (i.e., absolute accuracy), regardless of whether that misalignment is positively (overconfidence) or negatively (underconfidence) biased. The results of a one-way between-groups ANOVA found a statistically significant difference in mean calibration across accuracy quartile levels, Welch's $F(3,207) = 8.81, p < .001, \eta^2 = 0.11^3$. A post-hoc Games-Howell test comparing calibration among the four accuracy groups showed that mean calibration for the lowest performing group was significantly different from the middle two groups (both p 's $\leq .01$) but was not significantly different from the highest performing group ($p = .220$) indicating that low and higher performers were miscalibrated to a roughly equivalent degree (see Figure 2B).

Finally, though participants' self-reported estimates of their *general bullshit detection ability* appeared to followed a similar pattern to their task-specific overestimation when plotted (see Figure 2C), mean ratings among the four quartiles for this variable were not significantly different, $F(3,207) = 1.98, p = .118$. As this was a percentile measure, this indicates that most participants believe they have an above average ability to detect bullshit in general, regardless of actual ability ($M = 67.55, SD = 18.40$).

³ We report Welch's F -test here because Levene's test for equality of variances was violated, $p < .001$.

Figure 2. Graph A (left) represents mean differences bullshit detection overestimation by accuracy quartile. Graph B (middle) represents mean calibration index by accuracy quartile. Graph C represents mean self-reported “general bullshit detection ability” by accuracy quartile. 1 = lowest quartile; 4 = highest quartile. Error bars represent 95% confidence intervals.

Overplacement

We next conducted a one-way ANOVA comparing mean levels of overplacement⁴ across all accuracy quartiles, $F(3,207) = 112.97, p < .001, \eta^2 = 0.62$. A post-hoc Tukey HSD test revealed significant mean differences in overplacement among all groups (all p 's $< .001$). The degree to which participants overplaced their own performance relative to the average performance of others followed a similar pattern to their overestimation results (see Figure 3). Specifically, participants in the first (bottom) quartile ($M = 26.95, SD = 15.21$) overplaced their relative performance to a greater degree than those in the second quartile ($M = 7.48, SD = 14.00$) while those in the third ($M = -11.20, SD = 15.75$) and fourth ($M = -22.49, SD = 11.80$) quartiles significantly underplaced their performance. These findings suggest that people who are the most highly receptive to bullshit could be argued to have a “bullshit blind spot” in that they are not

⁴ Recall that overplacement is a measure of bias rather than a difference score. Scores greater than 0 reflect the degree to which a person's belief that their performance is better than the average person is unwarranted while scores below 0 represent the degree to which a person's belief that their performance is worse than the average person is unwarranted.

only poor at detecting bullshit, but they also believe they are better at it (i.e., less biased) than the average person (Moore & Healy, 2008; Moore & Schatz, 2017; Pronin et al., 2002; Zell et al., 2020).

Moreover, this specific type of “better-than-average” bias did not extend to individuals with superior bullshit detection ability, despite them being similarly unaware of their own detection performance relative to others. Indeed, contrary to the overplacement displayed by those in the lowest quartile, those in the highest quartile instead demonstrated relatively large “worse than average” beliefs, underestimating their performance relative to that of other participants (arguably a “blindsight” contrasting the low performers’ “blind spot”). Overall, these results suggest that, in terms of one’s bullshit detection ability relative to others, whether a person is very skilled or very unskilled they are likely unaware (e.g. Burson et al., 2006).

Figure 3. Mean overplacement by accuracy quartile. 1 = lowest quartile; 4 = highest quartile. Error bars represent 95% confidence intervals.

Previous work has often attributed greater susceptibility to misinformation to a lack of reflective engagement (e.g., Pennycook et al., 2015a). Therefore, it is plausible that some people are better able to detect bullshit because they engage in more effortful, deliberative processing

when evaluating it compared to those who are less able to detect misinformation. As such, they may also perceive that they are at a higher risk of making errors when evaluating ambiguous bullshit information which, in turn, undermines their confidence to a greater degree than likely would be expected when evaluating less ambiguous (mis)information (Dunn et al., 2019; Maksimenko et al., 2020). In Study 2, we explore this idea by examining the frequency with which people report using intuitive versus reflective thinking processes when evaluating pseudo-profound and intentionally-profound statements.

Study 2

Study 1 provided valuable insight into some metacognitive processes associated with bullshit detection. For Study 2, we sought to more deeply examine the metacognitive and phenomenological experience of bullshit detection at the individual level. Past research has found that bullshit receptivity is negatively associated with measures of analytic thinking, including the Cognitive Reflection Test (CRT) which is often claimed to be a measure of engagement in reflective thinking (e.g., Frederick, 2005; Littrell et al., 2020; Pennycook et al., 2015a). According to this model, higher bullshit receptivity is strongly driven by a lack of engagement in critical, reflective thinking processes. However, a growing body of work has provided strong evidence that a person's score on the CRT may not be a reliable indicator of state or trait engagement in reflective thinking per se (e.g., Blacksmith et al., 2019; Patel et al., 2019). Instead, higher scores on the CRT may sometimes result from “smart intuition” – an instance of a highly accurate, context-specific matching heuristic – rather than cognitive reflection (Bago & De Neys, 2019; Ghasemi et al., 2022; Raoelison et al., 2020). This may help to explain the finding of low performers on the CRT overreporting that they were “dispositionally analytic” thinkers (Pennycook et al., 2017). Moreover, even individuals with

greater levels of relevant knowledge and cognitive sophistication can commit errors in reasoning which leave them more vulnerable to various types of biases (Stahl & Prooijen, 2018; Stanovich, 2018; West et al., 2012). This suggests that models which attribute susceptibility to misleading information primarily to a lack of cognitive reflection may be insufficient, as they do not account for the critical roles that smart intuition and error-prone rationalization may play in bullshit detection. Additionally, we wanted to assess the degree to which people possess sufficient insight into their own thinking processes when evaluating misleading information, given that Dunning-Kruger-like effects are often attributed to metacognitive deficits. Thus, we were curious whether such presumed deficits might extend to self-perceptions of one's own analytic processes when evaluating bullshit.

To explore this idea, we introduced a measure intended to assess participants' perceptions of the thinking processes (i.e., intuition versus reflection) they utilize when evaluating potentially misleading information. We did this to examine whether participants' perceived engagement in specific thinking processes align with objective performance in ways that are theoretically-consistent with past research (e.g., Pennycook et al., 2015b). We also recorded time spent evaluating each statement (i.e., reaction times) as a more objective measure to corroborate participants' self-reports. These new measures replaced our overplacement and self-reported general bullshit detection ability variables from Study 1.

Method

Participants

Based on our goal of achieving at least .80 power to detect an effect of Cohen's $f = .25$ at $\alpha = .05$ (g*power; Faul, Erdfelder, Buchner, & Lang, 2009), we used CloudResearch to recruit

229 adult participants from the United States and Canada from Amazon's MTurk online participant pool (Litman, Robinson, & Abberbock, 2016). As with Study 1, only those who had completed a minimum of 500 surveys, had at least a 99% MTurk HIT approval rating, and passed the two "botcha" eligibility questions were eligible to participate. The botchas prevented 39 potential fake responders from accessing our survey. Of the 229 participants who were eligible for the survey, data for 28 were removed for failing an attention check, leaving us with data for 201 participants to analyze (112 males, 88 females, 1 prefer not to answer, $M_{\text{age}} = 38.55$, $SD_{\text{age}} = 12.35$, Bachelor's degree or higher = 60.9%). Participants were paid \$2.00 USD for the roughly 15-minute survey.

Procedure and Materials

Informed consent and online survey presentation procedures were identical to Study 1 (participants also completed measures of bullshitting frequency and verbal intelligence for reasons peripheral to the focus of the current paper, so they will not be discussed here). As with Study 1, participants completed the bullshit detection task and the overestimation item. Additionally, we created a measure intended to capture a person's metacognitive experience or phenomenology of bullshit detection. While completing the bullshit detection task, participants were asked to self-report whether they used intuition or reflection when evaluating each item, which provided a measure of perceived "thinking process." After rating each pseudo-profound or motivational item, participants were taken to a new screen where they read the following instructions: *"Please select the response below that best describes how you decided on this rating."* They were then asked to select one of the following two options based on how they had rated the previous item:

- 1) *“I read it and immediately knew it was profound (NOT profound).”*
- 2) *“I had to think about it for at least a few seconds before determining that it was profound (NOT profound).”*

Items were coded as “intuitive” (0) or “reflective” (1). We used Qualtrics skip-logic function so that the wording of the thinking process items would correspond to the rating the participant had given the previous item (i.e., profound or not profound). That is, if a participant rated an item as “profound,” the thinking process options on the next screen were worded to correspond with that rating (e.g., “I read it and immediately knew it was profound”). The first item (“I read it and immediately knew...”) was intended to indicate the use of an intuitive thinking process when evaluating the previous pseudo-profound or motivational statement. The second item (“I had to think about it for at least a few seconds...”) was meant to indicate that a more reflective process had been used.

Responses to each thinking process item were summed for pseudo-profound statements and motivational quotes separately, while a combined score was also calculated. Higher “thinking process” scores are intended to represent a greater perceived reliance on cognitive reflection (and less on intuition) when evaluating items while lower scores are interpreted as indicating less perceived reliance on reflective processes (thus more on intuition). To corroborate these self-reported thinking processes with a more objective measure, completion times were also recorded during the evaluation of each pseudo-profound or motivational statement using Qualtrics’ “question timing” feature. Average completion/evaluation times were calculated for pseudo-profound (bullshit) items, motivational (non-bullshit) items, and across all items.

Results and discussion

We found no significant associations between any of our main variables of interest and self-reported thinking process scores for pseudo-profound and motivational statements separately, therefore we collapsed the self-reported thinking process scores across all statement types. We then compared these self-reported thinking process scores to the objectively measured evaluation (response) times recorded by Qualtrics. The response time (RT) distributions for both pseudo-profound ($z(\text{skew}) = 31.31$) and motivational statements ($z(\text{skew}) = 19.47$) were significantly positively skewed, so they were log-transformed and collapsed across all items to align with our thinking process variable. Though “self-reported thinking process” is admittedly a somewhat “noisy” measure, scores for this variable were positively correlated with objectively-measured RTs, $r(200) = .16$, $p < .02$. However, given that this association is small and behavioural measures of complex psychological phenomena (e.g., thinking processes) have been shown to exhibit poor reliability and are often relatively weak predictors of real-world outcomes, we urge caution against overinterpreting this relation (c.f., Dang et al., 2020; Eisenberg et al., 2019).

Bullshit detection overestimation

We first sought to replicate our findings from Study 1 related to overestimation in bullshit detection. We split our sample into four groups based on bullshit detection accuracy quartiles and conducted a one-way ANOVA to compare differences in bullshit detection overestimation across quartile levels (see Figure 3), finding a statistically significant main effect, $F(3,197) = 53.21$, $p < .001$, $\eta^2 = .45$. A post-hoc Tukey HSD test comparing overestimation among the four groups showed significant mean differences between three out of the four groups ($p < .01$). The mean difference between the 3rd and 4th quartiles ($M = 9.28$), was not significant ($p = .07$).

Replicating findings from Study 1, results revealed that those who were the worst at detecting bullshit were the most overconfident in their performance while those who were the best at detecting bullshit were the most underconfident. Also similar to Study 1 results, the lowest performers were positively miscalibrated (i.e., overconfident) to a similar degree that the highest performers were negatively miscalibrated (i.e., underconfident). Specifically, participants who scored in the bottom quartile, ($M = 41.20$, $SD = 8.24$) overestimated their bullshit detection accuracy by 25.91% on average while those in the top quartile ($M = 85.70$, $SD = 5.63$) underestimated their detection accuracy by -19.12%. Finally, second quartile participants ($M = 61.36$, $SD = 3.53$) overestimated their accuracy by an average of only 4.35% and third quartile participants ($M = 72.38$, $SD = 2.53$) underestimated their accuracy by an average of -9.31%.

Figure 4. Graph A (left) represents mean differences bullshit detection overestimation by accuracy quartile. Graph B (right) represents mean calibration by accuracy quartile. 1 = lowest quartile; 4 = highest quartile. Error bars represent 95% confidence intervals.

Metacognitive experience of bullshit detection

As the main goal of Study 2 was to examine the thinking processes used to evaluate bullshit, we next conducted a series of exploratory analyses examining potential patterns in self-reported thinking processes used by participants when evaluating the bullshit detection task

items. Participants rated each statement (10 pseudo-profound, 10 motivational) according to which process (i.e., intuition or reflection) they used to determine whether it was profound. From this, we were able to assess the proportion of participants who self-reported utilizing intuition and the proportion who self-reported engaging in reflection for each bullshit and non-bullshit statements separately.

Our data collection was organized such that we were able to separate intuitive/reflective responses for items which were rated “profound” from those items which were rated “not profound” by each participant. However, overall, we found no differences in patterns of self-reported thinking processes for statements rated “profound” versus “not profound” (i.e., there was no evidence that a specific thinking process was dominant overall for either response type). Separate chi-square tests of independence were conducted for each of the 10 pseudo-profound statements and the 10 motivational quotes to examine whether there were significant differences in the thinking process participants reported using when evaluating each item (results for all 20 tests are listed in the Supplementary Materials). There were no significant differences in thinking process between most of the pseudo-profound items ($p > .10$). That is, for 7 out of the 10 pseudo-profound statements, roughly half of participants self-reported that they evaluated the item using intuition while the other half reported using reflection. However, we did find a significant difference ($p < .01$) in thinking process for 8 out of the 10 motivational statements such that a significant majority of participants self-reported evaluating those 8 items using intuition rather than reflection.

When directly comparing frequency of responses across all pseudo-profound statements, 51% of participants self-reported that they used intuition to evaluate pseudo-profound items versus 49% who self-reported using reflection. Across all motivational quotes, 64% reported that

they evaluated these items using intuition versus 36% who reported they used reflection. To illustrate the overall patterns, we provide results in Table 2 from four sample items as well as an overall comparison of results for pseudo-profound statements versus motivational quotes. Chi-square results for all 20 bullshit detection task items can be found in the Supplementary Materials for interested readers.

Table 2. Examples of intuition versus reflection use in BS detection

	<i>Item</i>	<i>Process</i>	<i>Count</i>	<i>Proportion</i>	χ^2	<i>Sig.</i>	<i>Avg. time</i>
1	Wholeness quiets infinite phenomena. (BS)	Intuition	103	.51	0.12	.72	6.95 s
		Reflection	98	.49			11.15 s
2	This life is nothing short of an unfolding lightning bolt of enlightened growth. (BS)	Intuition	102	.51	0.04	.83	9.00 s
		Reflection	99	.49			10.93 s
3	A wet person does not fear the rain. (non-BS)	Intuition	131	.66	19.22	<.001	6.13 s
		Reflection	69	.34			12.24 s
4	A river cuts through a rock not because of its power but its persistence. (non-BS)	Intuition	155	.77	59.11	<.001	6.61 s
		Reflection	46	.23			12.92 s
<i>Overall</i>		<i>Process</i>	<i>% of participants reported using</i>		<i>Average rating time per process</i>		
Pseudo-profound BS items (bullshit)		Intuition	50.6%		8.93 s		
		Reflection	49.4%		13.47 s		
Motivational quotes (non-bullshit)		Intuition	64.4%		7.43 s		
		Reflection	35.6%		12.60 s		

Note: Rating times reported in seconds.

General discussion

Across two studies ($N = 412$), we found consistent evidence that people who were least able – and those who were most able – to distinguish bullshit from non-bullshit were metacognitively miscalibrated with respect to their performance on a bullshit detection task.

Importantly, greater receptivity to bullshit was not adequately explained by self-reported thinking processes. That is, while most participants reported using intuitive processes to evaluate motivational statements (i.e., non-bullshit), no clear trend emerged in self-reported thinking process for pseudo-profound statements (i.e., bullshit). Indeed, participants were evenly split in how they evaluated pseudo-profound statements, with approximately half self-reporting that they used intuition while the other half reported using reflection. In the following, we expand on these results and their implications while offering suggestions for future avenues of research.

Of blind spots and blindsight

Past research on overconfidence effects in areas such as fake news detection (Lyons et al., 2021) and analytic thinking ability (Pennycook et al., 2017) has found that performers in the bottom quartile are generally miscalibrated to a larger degree than those in the top quartile. Moreover, this pattern – that the confidence-accuracy gap is measurably larger for the least accurate performers compared to the most accurate – is fairly consistent, having been found across a wide range of tasks in various domains (e.g., Kruger & Dunning, 1999; Jansen et al., 2021; Pennycook et al., 2017; Schlosser et al., 2013). The results reported here somewhat diverge from this pattern in that the degree to which the highest achievers were miscalibrated is arguably more pronounced than that found in past studies. Indeed, those who were especially poor at detecting bullshit and those who were especially good at it were nearly equally miscalibrated. This pattern also emerged for overplacement and underplacement.

Specifically, individuals who were especially poor at bullshit detection were also overconfident in their ability (both general and task-specific) to distinguish bullshit from non-bullshit, suggesting that highly bullshit receptive people are largely unaware of their own susceptibility to misleading information. Moreover, the lowest performers also believed that they

were better at detecting BS than the average person. This unfortunate combination of an obliviousness to one's lack of ability coupled with a perceived superiority compared to others might be reasonably construed as an example of a "bias blind spot" or, in this case, a "*bullshit blind spot*" (c.f. Pronin et al, 2002). Conversely, those in the highest bullshit detection accuracy quartile both underestimated and underplaced their performance to a similar degree as those in the lowest quartile overestimated and overplaced their own abilities. In a similar fashion to the bias blind spot, this double disadvantage of underestimating and underplacing one's superior abilities might be said to represent a form of "*bias blindsight*"⁵.

Taken together, these results suggest that whether a person is very skilled or very unskilled at detecting bullshit, they are likely ignorant to whether they are good or bad at it. Moreover, though research exists on instances of bias blind spots, the extent to which a *bias blindsight* might manifest in non-bullshit contexts is currently unclear and represents an interesting question for future research. Critically, however, given the potential harms associated with falling for misleading information, the bias *blind spot* of highly bullshit-receptive people is potentially more deleterious in the aggregate than the bias *blindsight* of those who are highly bullshit-resistant, which helps explain why receptivity to misleading information has received more research attention thus far.

The role of ambiguity in bullshit detection

One possible explanation for the results reported here is that since pseudo-profound bullshit statements are constructed using random, esoteric buzzwords with no intention to communicate a meaningful message, they represent a crucially more ambiguous type of

⁵ "Blindsight" refers to the phenomenon in which people with certain forms of blindness respond to visual stimuli which they are unable to consciously see.

information than the fake news headlines and analytic thinking tasks used in similar previous studies (e.g., Lyons et al., 2021; Pennycook et al., 2017). Although people often engage in more cognitive effort when processing ambiguous information, they often commit more errors than when processing information with less ambiguity (Dunn et al., 2019; Maksimenko et al., 2020). Recent work has shown that some individuals who are especially receptive to bullshit (e.g., persuasive bullshitters) appear metacognitively insensitive to linguistic cues that might help them better distinguish information which is superficially impressive but semantically meaningless from information that is intentionally profound/truthful and semantically meaningful (Littrell et al., 2021b). That is, it appears that at some level, many highly bullshit receptive individuals are less able to detect conflict in misleading information (or at least less able to bring such conflict detection to the level of conscious awareness), the difficulty of which may be exacerbated when processing highly ambiguous information (De Neys et al., 2010; De Neys, 2012; Dunn et al., 2019; Maksimenko et al., 2020; Pennycook et al., 2015b).

Furthermore, people who are *misinformed* are significantly more resistant to updating their beliefs compared to those who are simply *uninformed*, and those with the most inaccurate beliefs are not only more likely to avoid sources of corrective information, but more likely to discount the accuracy of those sources (Li & Wagner, 2020; Sheldon et al., 2014). These findings suggest that overconfidence in one's ability to detect misleading or epistemically-suspect information (e.g., bullshit) could include an increased resistance to engaging with fact-checking resources or other corrective interventions (Lyons et al., 2021; Sheldon et al., 2014) and less willingness to revise one's beliefs in the face of new information (Li & Wagner, 2020; Walter et al., 2020). These types of negative downstream effects would be less likely to occur for those who are highly skilled, yet underconfident, in detecting misleading information.

Dual process theories and bullshit detection

Our findings provide partial support for both an intuitive and a reflective account of bullshit detection, as nearly 50% of participants indicated that they evaluated pseudo-profound statements intuitively rather than reflectively (which also corresponded with response times), and this was the case for both high and low bullshit detection performance. Though this may appear somewhat contradictory, as well as in conflict with proposed miserly processing accounts of bullshit receptivity (e.g., Pennycook et al., 2015a; Pennycook & Rand, 2020), our results align quite well with the three-stage dual process model of analytic engagement proposed by Pennycook, Fugelsang, and Koehler (2015). In their model, human reasoning progresses from an initial intuitive response (stage 1) to a conflict monitoring process (stage 2). If no conflict is detected at this stage, the person will revert to their initial intuitive response. However, if conflict is detected, reasoning proceeds to one of two types of deliberative engagement (Stage 3): 1) *cognitive decoupling*, where one engages critical, analytic processes to override an incorrect, intuitive response, or; 2) *rationalization*, which describes a more biased and error-prone form of reflection in which a person generates justifications in support of an initial incorrect response (i.e., talking oneself into the wrong answer).

The majority of participants who correctly identified the intentionally-profound motivational quotes as “profound,” reported using intuitive processes when evaluating these statements. Additionally, approximately half of participants who correctly identified the meaningless, pseudo-profound statements as “not profound” reported using intuition when making these judgments. These results align with theoretical accounts of Stage 1 “smart intuitors” – people who can quickly and accurately evaluate certain types of information using a type of experience-based matching heuristic (Bago & De Neys, 2019; Ghasemi et al., 2022; Raoelison et al., 2020; Stanovich, 2018). Conversely, roughly 50% of participants who rated

bullshit statements as “profound” reported using intuition to evaluate these statements, suggesting that they lack the requisite knowledge structures and analytic thinking framework to detect bullshit at the instantiation of intuition (Stage 1) or conflict monitoring (Stage 2). Moreover, the fact that participants reported using reflective thinking for approximately half of all “profound” and half of all “not profound” ratings for pseudo-profound statements appears to support the proposed Stage 3 reflective processes of *cognitive decoupling* and *rationalization*. That is, those who evaluated bullshit statements as “not profound” likely engaged in cognitive decoupling while those who rated such statements as “profound,” most likely utilized sub-optimal rationalization.

Given the consistent negative associations between bullshit receptivity and analytic thinking measures found in past research, some have suggested that people can decrease their susceptibility to misleading information by engaging in more reflective, analytic thinking in general (Pennycook et al., 2020). While this may be broadly helpful, the results presented here suggest that some highly bullshit receptive individuals engage in ineffective, error-prone reflective processes. Therefore, future research would benefit from a deeper examination of the reflective processes used when evaluating misleading information to uncover the sources of these reasoning errors. This could involve focusing on interventions that encourage more specific types of reflective thinking strategies, rather than simply promoting reflection, per se. Indeed, recent work utilizing interventions such as accuracy nudges and educational video games that seek to increase analytic engagement have shown some success in decreasing people’s susceptibility to falling for misinformation and their willingness to share it online, illustrating the benefit of novel reflective approaches (Roozenbeek & van der Linden, 2019; Pennycook et al., 2021). However, as evidenced in the present findings, future research would also benefit from

focusing on the ways in which overconfidence in one's detection abilities might interact with or otherwise adversely impact the effectiveness of these types of interventions.

Criticisms of Dunning-Kruger effects

Although miscalibration effects similar to those presented here (often referred to as *Dunning Kruger effects*) have been found in a variety of domains including face perception (Zhou & Jenkins, 2020), anti-vaccine attitudes (Motta et al., 2018), opposition to genetically-modified foods (Fernbach et al., 2019), fake news detection (Lyons et al., 2021), and analytic thinking (Pennycook et al., 2017), they should not be assumed to be a foregone conclusion for all phenomena. Indeed, there is evidence to suggest that such effects may not necessarily generalize across cognitive tasks, even within the same domain (Gignac, 2022; Hofer et al., 2021). Moreover, these effects do not always generalize across different types of overconfidence (Moore & Healy, 2008). Research has found that three common forms of overconfidence (i.e., overestimation, overplacement, and overprecision) are sometimes only marginally related depending on the task (Moore & Healy, 2008; Moore & Schatz, 2017) and sometimes differ in the ways in which they relate to other variables of interest (e.g., Macenczak et al., 2016). Moreover, a Dunning-Kruger-like effect can exist for one type of overconfidence measure yet not for others. For example, in a study of U.S. Department of Energy employees' knowledge of nuclear weapons science and technology, Sanchez and Speed (2020) found a classic Dunning-Kruger-like pattern for overplacement but not for overestimation or overprecision. Additionally, experts have been found to display higher overestimation and overplacement than novices in the areas of financial literacy and stock market predictions, rather than the underestimation that would be more typical of DKE patterns found in general population samples (Kang & Kim, 2021; Sanchez & Dunning, 2018).

Additionally, though the so-called Dunning Kruger effect is one of the more highly replicable and robust effects in psychology (e.g., Jansen et al., 2021; Lyons et al., 2021; Pennycook et al., 2017), it is also not without its critics. Indeed, there has been some debate as to whether these patterns of results reflect individual differences in metacognition (e.g., Dunning et al., 2003; Kruger & Dunning, 1999; McIntosh et al., 2019), the “better-than-average” heuristic (Krueger & Mueller, 2002), or biases motivated toward self-enhancement (e.g., Anderson et al., 2012; Roy & Leirsch, 2013; Zell et al., 2020). While we find each of these accounts compelling in their own ways, we believe the patterns in our data, as well as the bulk of the more current relevant literature, support an interpretation that our results are best explained by factors related to metacognition and, to a lesser degree, self-enhancement.

It should also be noted that some have argued that Dunning-Kruger-like effects are simply statistical artifacts of regression to the mean paired with the “better-than-average heuristic” (Gignac & Zajenkowski, 2020; Krueger & Mueller, 2002). To their credit, Kruger and Dunning (2002) have responded directly to these criticisms, arguing that a re-analysis of their original data did not support a regression-to-the-mean explanation. Subsequent results from other researchers, which have included metaanalyses and replications using data from more real-world situations, have provided further support that a metacognitive account, rather than regression to the mean, is the more appropriate interpretation of the data (e.g., Burson et al., 2006; Coutinho et al., 2021; Ehrlinger et al., 2008; Jansen et al., 2021; Pennycook et al., 2017; Schlösser et al., 2013).

Crucially, perhaps the strongest evidence supporting a metacognitive account for Dunning-Kruger-like effects comes from more recent studies, such as that by Jansen and colleagues (2021). Combining both a computational modelling approach with a large-scale direct

replication ($N = 7068$) of Kruger's and Dunning's (1999) original study, Jansen and colleagues established that so-called Dunning-Kruger effects were best explained by a rational model which assumes significantly reduced sensitivity to errors among poor performers (Jansen et al., 2021; Mazar & Fleming, 2021). Indeed, these results are supported by recent evidence finding that, for performance on the CRT, overconfidence in the lowest performers was predicted by errors in intuitive processes which were not present for the highest performers (Coutinho et al., 2021). This account also aligns with recent work showing that people scoring higher in bullshit receptivity appear to be largely insensitive to linguistic differences in bullshit and non-bullshit, leading them to misjudge the superficial profoundness of a message as roughly equivalent to its inherent profoundness (Littrell et al., 2021b).

Finally, data from recent EEG studies has revealed clear, significant differences between underestimators and overestimators at both the behavioural level (e.g., response time measures) and neurophysiological level (e.g., event-related potential effects) for Dunning-Kruger memory tasks (Muller et al., 2021). These differences in brain activity suggest that different cognitive processes are involved in under- and over-estimation and are not adequately explained by regression-to-the mean accounts of Dunning-Kruger-like effects. (Muller et al., 2021).

Limitations

Though there are various types/domains of bullshit (e.g., scientific, organizational, political, fake news, etc.), the present study focused only on the detection of pseudo-profound bullshit, which is composed of abstruse words and phrases that are arguably more pseudo-metaphysical than pseudo-factual. As such, the present results may not generalize to types/domains of bullshit that rely on cues more evocative of factual statements than metaphysical profundities.

Additionally, the present study used English-speaking participants from Canada and the United States. Therefore, we urge caution in generalizing these results to non-Western cultures and populations whose primary language is not English. Indeed, past research has uncovered cultural differences in detecting deception (Taylor, et al., 2015) and it is reasonable to assume that there exists in other cultures a wide variety of conceptualizations of bullshit and strategies for detecting it. Finally, this study looked only at cognitive and metacognitive factors associated with bullshit detection. There are also important situational, emotional, and social factors that may exert measurable influences on bullshit detection that would be well-served by an approach that incorporates more personality-based and social psychological perspectives (e.g., Forgas & East, 2008; Petrocelli, 2021).

Conclusion

The results reported here provide empirical evidence that the people most vulnerable to falling for at least one type of bullshit are not only unaware of their susceptibility but also believe themselves to be superior to others at detecting it (i.e., bullshit blind spot). Conversely, people who are the best at detecting bullshit are unaware of their superior ability and believe themselves to be slightly worse than others at detecting it (i.e., bullshit blindsight). Additionally, both intuitive and reflective thinking processes were found to play critical roles in successful and unsuccessful bullshit detection, underscoring the need for researchers to more closely consider the ways in which both types of processes can reliably lead to optimal (and sub-optimal) reasoning outcomes as well as continue to investigate individual differences which may influence the success or failure of each reasoning process in different contexts. Indeed, though the intuition of some participants was accurate and reliable for distinguishing bullshit from non-bullshit, the reflective processes used for bullshit detection by others was prone to errors and

bias. Put another way, just as there are “smart intuitors,” there are also “poor reflectors.” Thus, a more fruitful path in the fight against misinformation may lie not simply in encouraging people to reflect, but in showing them better ways to reflect.

References

- Anderson, C., Brion, S., Moore, D. A., & Kennedy, J. A. (2012). A status-enhancement account of overconfidence. *Journal of Personality and Social Psychology*, 103(4), 718–735.
<https://doi.org/10.1037/a0029395>
- Bago, B., & De Neys, W. (2019). The Smart System 1: Evidence for the intuitive nature of correct responding on the bat-and-ball problem. *Thinking & Reasoning*, 25(3), 257–299.
<https://doi.org/10.1080/13546783.2018.1507949>
- Bago, B., Rand, D. G., & Pennycook, G. (2020). Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines. *Journal of Experimental Psychology: General*, 149(8), 1608-1613. <http://dx.doi.org/10.1037/xge0000729>
- Blacksmith, N., Yang, Y., Behrend, T. S., & Ruark, G. A. (2019). Assessing the validity of inferences from scores on the cognitive reflection test. *Journal of Behavioral Decision Making*, 32, 599– 612. <https://doi.org/10.1002/bdm.2133>
- Burson, K. A., Larrick, R. P., & Klayman, J. (2006). Skilled or unskilled, but still unaware of it: How perceptions of difficulty drive miscalibration in relative comparisons. *Journal of Personality and Social Psychology*, 90(1), 60-77.
<https://psycnet.apa.org/fulltext/2006-01254-004.html>
- Čavojová, V., Brezina, I., & Jurkovič, M. (2020). Expanding the bullshit research out of pseudo-transcendental domain. *Current Psychology*, Advance online publication, 1-10.
<https://doi.org/10.1007/s12144-020-00617-3>

- 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>
- Chmielewski, M., & Kucker, S. C. (2020) An MTurk Crisis? Shifts in Data Quality and the Impact on Study Results. *Social Psychological and Personality Science*, 11(4):464-473. doi:10.1177/1948550619875149
- 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>
- Coutinho, M., Thomas, J., Alsuwaidi, A., & Couchman, J. J. (2021). Dunning-Kruger Effect: Intuitive Errors Predict Overconfidence on the Cognitive Reflection Test. *Frontiers in Psychology*, 12, 603225. <https://doi.org/10.3389/fpsyg.2021.603225>
- Dang, J., King, K. M., & Inzlicht, M. (2020). Why are self-report and behavioral measures weakly correlated? *Trends in Cognitive Sciences*, 24(4), 267-269. <https://doi.org/10.1016/j.tics.2020.01.007>
- De keersmaecker, J., & Roets, A. (2017). ‘Fake news’: Incorrect, but hard to correct. The role of cognitive ability on the impact of false information on social impressions. *Intelligence*, 65, 107-110. <https://doi.org/10.1016/j.intell.2017.10.005>
- De Neys, W. (2012). Bias and conflict: A case for logical intuitions. *Perspectives on Psychological Science*, 7(1), 28–38. <https://doi.org/10.1177%2F1745691611429354>
- De Neys, W., Moyens, E., & Vansteenwegen, D. (2010). Feeling we’re biased: Autonomic arousal and reasoning conflict. *Cognitive, Affective, & Behavioral Neuroscience*, 10(2), 208–216. <https://doi.org/10.3758/CABN.10.2.208>

Dunn, T. L., Inzlicht, M., & Risko, E. F. (2019). Anticipating cognitive effort: roles of perceived error-likelihood and time demands. *Psychological Research*, 83(5), 1033-1056.

<https://doi.org/10.1007/s00426-017-0943-x>

Dunning, D., Johnson, K., Ehrlinger, J., & Kruger, J. (2003). Why people fail to recognize their own incompetence. *Current Directions in Psychological Science*, 12(3), 83–87.

<https://doi.org/10.1111/1467-8721.01235>

Ehrlinger, J., Johnson, K., Banner, M., Dunning, D. & Kruger, J. (2008). Why the unskilled are unaware: further explorations of (absent) self-insight among the incompetent.

Organizational Behavior and Human Decision Processes, 105(1), 98–121.

<https://doi.org/10.1016/j.obhdp.2007.05.002>

Eisenberg, I. W., Bissett, P. G., Zeynep Enkavi, A., Li, J., MacKinnon, D. P., Marsch, L. A., & Poldrack, R. A. (2019). Uncovering the structure of self-regulation through data-driven ontology discovery. *Nature Communications*, 10(1), 1-13.

<https://doi.org/10.1038/s41467-019-10301-1>

Evans, A. M., Sleegers, W., & Mlakar, Ž. (2020). 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(4), 1149–1160. <https://doi.org/10.3758/BRM.41.4.1149>

Fernbach, P. M., Light, N., Scott, S. E., Inbar, Y., & Rozin, P. (2019). Extreme opponents of genetically modified foods know the least but think they know the most. *Nature Human*

Behaviour, 3, 251-256. <https://doi.org/10.1038/s41562-018-0520-3>

- Forgas, J. P., & East, R. (2008). On being happy and gullible: Mood effects on skepticism and the detection of deception. *Journal of Experimental Social Psychology*, 44(5), 1362–1367. <https://doi.org/10.1016/j.jesp.2008.04.010>
- Frankfurt, Harry G. 2006. *On truth*. New York: Alfred A. Knopf.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4), 25-42. doi: 10.1257/089533005775196732
- Ghasemi, O., Handley, S., Howarth, S., Newman, I. R., & Thompson, V. A. (2022). Logical intuition is not really about logic. *Journal of Experimental Psychology: General*. Advance online publication. <https://doi.org/10.1037/xge0001179>
- Gignac, G. E. (2022). The association between objective and subjective financial literacy: Failure to observe the Dunning-Kruger effect. *Personality and Individual Differences*, 184, 111224. Advance online publication. <https://doi.org/10.1016/j.paid.2021.111224>
- Gignac, G. E., & Zajenkowski, M. (2020). The Dunning-Kruger effect is (mostly) a statistical artefact: Valid approaches to testing the hypothesis with individual differences data. *Intelligence*, 80, 101449. <https://doi.org/10.1016/j.intell.2020.101449>
- Hofer, G., Mraulak, V., Grinschgl, S., & Neubauer, A. (2021, November 24). Less intelligent and unaware? Accuracy and Dunning-Kruger effects for self-estimates of different aspects of intelligence. *PsyArxiv*, <https://doi.org/10.31234/osf.io/xavqh>
- Jansen, R. A., Rafferty, A. N., & Griffiths, T. L. (2021). A rational model of the Dunning-Kruger effect supports insensitivity to evidence in low performers. *Nature Human Behaviour*, 10.1038/s41562-021-01057-0. Advance online publication. <https://doi.org/10.1038/s41562-021-01057-0>

- Juul, J. L., & Ugander, J. (2021). Comparing information diffusion mechanisms by matching on cascade size. *Proceedings of the National Academy of Sciences*, 118(46), e2100786118. <https://doi.org/10.1073/pnas.2100786118>
- Kang, S. H., & Kim, J. (2021). The fragility of experts: A moderated-mediation model of expertise, expert identity threat, and overprecision. *Academy of Management Journal*, Advance online publication. <https://doi.org/10.5465/amj.2019.0899>
- Kienzler, M., Västfjäll, D., & Tinghög, G. (2022). Individual differences in susceptibility to financial bullshit. *Journal of Behavioral and Experimental Finance*, 34, 100655. <https://doi.org/10.1016/j.jbef.2022.100655>
- Krueger, J., & Mueller, R. A. (2002). Unskilled, unaware, or both? The better-than-average heuristic and statistical regression predict errors in estimates of own performance. *Journal of Personality and Social Psychology*, 82(2), 180-188. <https://doi.org/10.1037/0022-3514.82.2.180>
- Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, 77(6), 1121–1134. <https://doi.org/10.1037/0022-3514.77.6.1121>
- Kruger, J., & Dunning, D. (2002). Unskilled and unaware--but why? A reply to Krueger and Mueller (2002). *Journal of Personality and Social Psychology*, 82(2), 189–192. <https://doi.org/10.1037/0022-3514.82.2.189>
- Li, J., & Wagner, M. W. (2020). The value of not knowing: Partisan cue-taking and belief updating of the uninformed, the ambiguous, and the misinformed. *Journal of Communication*, 70(5), 646-669. <https://doi.org/10.1093/joc/jqaa022>

- Litman, L., Robinson, J., & Abberbock, T. (2016). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 49, 433-442. <https://doi.org/10.3758/s13428-016-0727-z>
- Littrell, S., Fugelsang, J. A., & Risko, E. F. (2020). Overconfidently underthinking: narcissism negatively predicts cognitive reflection, *Thinking & Reasoning*, 26(3), 352-380. <https://doi.org/10.1080/13546783.2019.1633404>
- Littrell, S., Risko, E. F., & Fugelsang, J. A. (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*, 60(4), 1484-1505. <https://doi.org/10.1111/bjso.12447>
- Lyons, B. A., Montgomery, J. M., Guess, A. M., Nyhan, B., & Reifler, J. (2021). Overconfidence in news judgments is associated with false news susceptibility. *Proceedings of the National Academy of Sciences*, 118(23), 1-10. <https://doi.org/10.1073/pnas.2019527118>
- Macenczak, L. A., Campbell, S., Henley, A. B., & Campbell, W. K. (2016). Direct and interactive effects of narcissism and power on overconfidence. *Personality and Individual Differences*, 91, 113-122. Retrieved from <https://doi.org/10.1016/j.paid.2015.11.053>

Maksimenko, V. A., Kuc, A., Frolov, N. S., Khramova, M. V., Pisarchik, A. N., & Hramov, A.

E. (2020). Dissociating cognitive processes during ambiguous information processing in perceptual decision-making. *Frontiers in Behavioral Neuroscience*, 14(95), 1-9.

<https://doi.org/10.3389/fnbeh.2020.00095>

Mazor, M., & Fleming, S. M. (2021). The Dunning-Kruger effect revisited. *Nature Human Behaviour*, 10.1038/s41562-021-01101-z. Advance online publication.

<https://doi.org/10.1038/s41562-021-01101-z>

McIntosh, R. D., Fowler, E. A., Lyu, T., & Della Sala, S. (2019). Wise up: Clarifying the role of metacognition in the Dunning-Kruger effect. *Journal of Experimental Psychology: General*, 148(11), 1882-1897.

<http://dx.doi.org/10.1037/xge0000579>

Merkle, C., & Weber, M. (2011). True overconfidence: The inability of rational information processing to account for apparent overconfidence. *Organizational Behavior and Human Decision Processes*, 116(2), 262–271.

<https://doi.org/10.1016/j.obhdp.2011.07.004>

Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*,

115(2), 502-517. <http://dx.doi.org/10.1037/0033-295X.115.2.502>

Moore, D. A., & Schatz, D. (2017). The three faces of overconfidence. *Social and Personality Psychology Compass*, 11(8), 1–12.

<https://doi.org/10.1111/spc3.12331>

Motta, M., Callaghan, T., & Sylvester, S. (2018). Knowing less but presuming more: Dunning-Kruger effects and the endorsement of anti-vaccine policy attitudes. *Social Science & Medicine*, 211, 274–281.

<https://doi.org/10.1016/j.socscimed.2018.06.032>

Muller, A., Sirianni, L. A., & Addante, R. J. (2021). Neural correlates of the Dunning–Kruger effect. *European Journal of Neuroscience*, 53(2), 460–484.

<https://doi.org/10.1111/ejn.14935>

- Ophir, Y., Sisso, I., Asterhan, C. S. C., Tikochinski, R., & Reichart, R. (2020). The Turker Blues: Hidden Factors Behind Increased Depression Rates Among Amazon's Mechanical Turkers. *Clinical Psychological Science*, 8(1), 65–83.
<https://doi.org/10.1177/2167702619865973>
- Patel, N., Baker, S. G., & Scherer, L. D. (2019). Evaluating the cognitive reflection test as a measure of intuition/reflection, numeracy, and insight problem solving, and the implications for understanding real-world judgments and beliefs. *Journal of Experimental Psychology: General*, 148(12), 2129–2153. <http://dx.doi.org/10.1037/xge0000592>
- Pennycook, G., Cheyne, J. A., Barr, N., Koehler, D. J., & Fugelsang, J. A. (2015a). 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., Fugelsang, J. A., & Koehler, D. J. (2015b). What makes us think? A three-stage dual-process model of analytic engagement. *Cognitive Psychology*, 80, 34–72.
<https://doi.org/10.1016/j.cogpsych.2015.05.001>
- Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A. A., Eckels, D., & Rand, D. G. (2021). Shifting attention to accuracy can reduce misinformation online. *Nature*, Advance online publication. <https://doi.org/10.1038/s41586-021-03344-2>
- 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>

Petrocelli, J. V. (2021). Politically oriented bullshit detection: Attitudinally conditional bullshit receptivity and bullshit sensitivity. *Group Processes & Intergroup Relations*,

1368430220987602. <https://doi.org/10.1177%2F1368430220987602>

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>

Raoelison, M., Thompson, V. A., & De Neys, W. (2020). The smart intuitor: Cognitive capacity predicts intuitive rather than deliberate thinking. *Cognition*, 204, 104381.

<https://doi.org/10.1016/j.cognition.2020.104381>

Roozenbeek, J., & van der Linden, S. (2019). Fake news game confers psychological resistance against online misinformation. *Palgrave Communications*, 5, 1-10.

<https://doi.org/10.1057/s41599-019-0279-9>

Roy, M. M., & Liersch, M. J. (2013). I am a better driver than you think: Examining self-enhancement for driving ability. *Journal of Applied Social Psychology*, 43(8),

1648-1659. <https://doi.org/10.1111/jasp.12117>

Salovich, N. A., & Rapp, D. N. (2021). Misinformed and unaware? Metacognition and the influence of inaccurate information. *Journal of Experimental Psychology: Learning,*

Memory, and Cognition, 47(4), 608–624. <https://doi.org/10.1037/xlm0000977>

- Sanchez, C., & Dunning, D. (2018). Overconfidence among beginners: Is a little learning a dangerous thing? *Journal of Personality and Social Psychology*, 114(1), 10–28.
<https://doi.org/10.1037/pspa0000102>
- Sanchez, Danielle J., & Speed, Ann. (2020). The impact of individual traits on domain task performance: Exploring the Dunning-Kruger Effect. United States Department of Energy.
<https://doi.org/10.2172/1716582>
- Schlösser, T., Dunning, D., Johnson, K. L., & Kruger, J. (2013). How unaware are the unskilled? Empirical tests of the “signal extraction” counter-explanation for the Dunning-Kruger Effect in self-Evaluation of performance. *Journal of Economic Psychology*, 39, 85-100. <http://dx.doi.org/10.1016/j.joep.2013.07.004>
- Schraw, G. (2009). A conceptual analysis of five measures of metacognitive monitoring. *Metacognition and Learning*, 4(1), 33–45. <https://doi.org/10.1007/s11409-008-9031-3>
- Scopelliti, I., Morewedge, C. K., McCormick, E., Min, H. L., Lebrecht, S., & Kassam, K. S. (2015). Bias blind spot: Structure, measurement, and consequences. *Management Science*, 61(10), 2468–2486. <https://doi.org/10.1287/mnsc.2014.2096>
- Sheldon, O. J., Dunning, D., & Ames, D. R. (2014). Emotionally unskilled, unaware, and uninterested in learning more: Reactions to feedback about deficits in emotional intelligence. *Journal of Applied Psychology*, 99(1), 125.
<https://doi.org/10.1037/a0034138>
- Simpson, S. (2019). *Fake News: A global epidemic, vast majority (86%) of online global citizens have been exposed to it*. Ipsos.
<https://www.ipsos.com/en-us/news-polls/cigi-fake-news-global-epidemic>

- Ståhl, T., & Prooijen, J.V. (2018). Epistemic rationality: Skepticism toward unfounded beliefs requires sufficient cognitive ability and motivation to be rational. *Personality and Individual Differences*, 122(1), 155-163. <https://doi.org/10.1016/j.paid.2017.10.026>
- Stanovich, K. E. (2018). Miserliness in human cognition: The interaction of detection, override and mindware. *Thinking & Reasoning*, 24(4), 423–444.
<https://doi.org/10.1080/13546783.2018.1459314>
- Taylor, P. J., Larner, S., Conchie, S. M., & van der Zee, S. (2015). Cross-cultural deception detection. In P. A. Granhag, A. Vrij, & B. Verschuere (Eds.), Wiley series in the psychology of crime, policing and law. *Detecting deception: Current challenges and cognitive approaches* (p. 175–201). Wiley-Blackwell.
- van Prooijen, J. W., Cohen Rodrigues, T., Bunzel, C., Georgescu, O., Komáromy, D., & Krouwel, A. P. (2022). Populist gullibility: Conspiracy theories, news credibility, bullshit receptivity, and paranormal belief. *Political Psychology*, Advance online publication.
<https://doi.org/10.1111/pops.12802>
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151. <https://doi.org/10.1126/science.aap9559>
- Walter, N., Cohen, J., Holbert, R. L., & Morag, Y. (2020). Fact-checking: A meta-analysis of what works and for whom. *Political Communication*, 37(3), 350-375.
<https://doi.org/10.1080/10584609.2019.1668894>
- West, R. F., Meserve, R. J., & Stanovich, K. E. (2012). Cognitive sophistication does not attenuate the bias blind spot. *Journal of Personality and Social Psychology*, 103(3), 506–519. <https://doi.org/10.1037/a0028857>

Zell, E., Strickhouser, J. E., Sedikides, C., & Alicke, M. D. (2020). The better-than-average effect in comparative self-evaluation: A comprehensive review and meta-analysis.

Psychological Bulletin, 146(2), 118–149. <https://doi.org/10.1037/bul0000218>

Zhou, X., Jenkins, R. (2020). Dunning–Kruger effects in face perception. *Cognition*,

203:104345. <https://doi.org/10.1016/j.cognition.2020.104345>