

Psychology of Belief | Talking to Deniers | Anne Frank's Betrayal | Havana Syndrome | Holy Honey | NASA and UFOs

Skeptical Inquirer

THE MAGAZINE FOR SCIENCE AND REASON

Vol. 47, No. 5 | September/October 2023

Persuasive Bullsh*tters



\$5.99 CAN/US



0 09281 02614 3

NEIL

Persuasive Bullshitters and the Insidious Bullshit Hypothesis

Short of lying, bullshitting involves communicating something with little to no regard for the truth, established knowledge, or genuine evidence. Although most people believe bullshit is relatively harmless, part of what makes bullshit compelling and appealing is that it often sounds beneficial, confirming, and credible. Four experiments in persuasion suggest that bullshit can have more influence than both lies and, in some cases, evidence-based frames of the very same information.

JOHN V. PETROCELLI

Claptrap, baloney, buffoonery, codswallop, empty-talk, garbage, hogwash, nonsense, poppycock, flapdoodle, flim-flam, flummery—and often better referenced as bullshit—our world appears to be full of it. Yet what exactly is bullshit? Bullshit is what emerges when people communicate with little to no regard for the truth, established knowledge, or genuine evidence (Frankfurt 1986; Petrocelli 2018; Petrocelli 2021a). Bullshitting behavior involves a broad array of rhetorical strategies that help us sound like we know what we’re talking about to impress others, persuade others, influence others, or explain things in an area in which our obligations to provide our opinions exceed our knowledge in that area (Cheyne and Pennycook 2016; Littrell et al. 2020; Petrocelli 2018; Petrocelli 2021a).

Bullshit is often abstract, lacking in sources and logic, or riddled with acronyms and business-speak jargon. Part of what makes bullshit compelling and appealing is that it often sounds beneficial, confirming, and credible—and thereby may be persuasive.

Much of the empirical research on bullshitting has dealt with the *when question*—that is, understanding *when* people are likely to bullshit us. Here, I offer a closer focus on the evidence that addresses part of the *why question*, that is, understanding the potential functions of bullshit. If bullshit is an undesirable communicative substance, *why* do people engage in the behavior so very often? Better understandings of when and why people bullshit us should place us in a much better position to detect and dispose of this unwanted social substance.

Bullshitting Isn’t Lying

Bullshit is not the same as lying (Frankfurt 1986). Although the liar and bullshitter are both deceptive in that they both *appear* concerned with the truth, only



MEBL

the liar is actually concerned with truth; the bullshitter is not. When someone lies to us, their objective is to get us to believe something they don't believe is true. If your colleague were to say something and he knows perfectly well it just isn't so, then he is lying. Whereas the liar doesn't believe what she says is true, the bullshitter has no idea whether or not it's true and, frankly, doesn't care. In fact, sometimes—just by chance or accident—the bullshitter says something that is true but even she wouldn't know it because she's not concerned about the truth, established knowledge, or evidence that may otherwise support her claims.

An important difference between lying and bullshitting involves the social reactions to these two forms of communication. When people lie to us, we're very unhappy with them, and they often pay great consequences for their lies. People tend to react to the lies with strong disdain and contempt, and there may be severe consequences for being caught in a lie (e.g., "She's a damn liar. That woman should be fired!"). In the case of bullshitting, the social consequences are much less severe (Petrocelli, Silverman, et al. 2023). The bullshitter may communicate the very same message as the liar, yet because she doesn't really know what she's talking about—and neither does she really care what the evidence suggests—we typically pass off her bullshit as a mild social offense (e.g., "Oh, she's *just* bullshitting again."). We assume the bullshit is harmless—but this is where we couldn't be more wrong.

Insidious Bullshit Hypothesis

Although bullshitting is often evaluated less negatively than lying (Petrocelli, Silverman et al. 2023), Harry Frankfurt (1986) proposed that bullshit is more damaging to society than lies. Not only by definition is there likely to be more bullshit than lies in the world, Frankfurt believed that pursuit of truth is of paramount importance—and although the liar's concern for truth is a corrupt one, it is less bad than the bullshitter's total disregard for truth.

Through my research in my Bullshit Studies Lab (at Wake Forest University), we've considered some additional possibilities that test this hypothesis. First, there are reasons to believe that bullshit may be more influential and persuasive than both evidence-based communications and lies. Second, there are also reasons to suspect that bullshit may interfere with memory more than lies and, thereby, have a greater impact on what people believe is true.

Persuasiveness of Bullshit

In our first experiment, we were interested in determining whether or not bullshit has a noteworthy influence on attitudes and beliefs when examined in light of a traditional persuasion procedure (Petrocelli 2021b). Over 500 college students were led to believe there was a new policy being proposed by university administrators that would require seniors to pass a comprehensive exam in their major area to graduate. Students were led to believe we were interested in surveying their thoughts and opinions about the proposed policy. Half

the students read persuasive arguments in favor of the exam policy that were either *strong* (e.g., Students who graduate from schools with comprehensive exams earn more in their starting salaries) or *weak* (e.g., Duke University is doing it). Half the students were also assigned to either a *bullshit frames condition* that prefaced the arguments with comments suggesting little to no interest in available and genuine evidence (e.g., "I believe there is some research on this issue, but I'm not really concerned with the evidence") or an *evidence-based frames condition* that prefaced the arguments with comments suggesting considerable interest in the relevant evidence (e.g., "I'm really concerned with the evidence concerning this issue").

Our analysis showed clear evidence of the traditional *Argument Quality effect*, but only within the evidence-based frame condition, which suggests a cueing of central route processing (whereby persuasion can occur only when the arguments themselves are compelling). However, the lack of a traditional *Argument Quality effect* within the bullshit frame condition suggests that bullshitting may cue what is usually referred to as peripheral route processing (whereby persuasion can occur when peripheral cues such as the number of arguments or attributes of the communicator appear compelling; see Figure 1).

From another angle, relative to evidenced-based frames, bullshit frames appeared to weaken the potency of strong arguments yet strengthen the potency of weak arguments. Such findings are quite remarkable given the arguments themselves are the very same—they only differ in what they signal about the intentions of the communicator as they pertain to their concern for truth and evidence.

Theoretically, there are two general routes to persuasion of which the route depends on the ability, opportunity, and mental resources one has at the time, as well as their general motivation to think about the content of the message. When people follow the *central route to persuasion*—because they have the ability, opportunity, and motivation to think about a persuasive communication—they tend to listen carefully and




Figure 1. Attitude and Thought Favorability means by Argument Quality and Argument Frame (Experiment 1); error bars represent +/-1 standard error.

think about the quality of the arguments presented. When people follow the *peripheral route to persuasion*—because they are either not motivated or lack the capacity to think about the arguments in a persuasive communication—they may be swayed by more superficial cues (e.g., source credibility, source attractiveness, nonverbal cues, number of arguments in the message) rather than their quality.

We conducted another experiment; this time, the attitude object was student feelings about an extended two-week Spring Break—an idea most students like (Petrocelli 2021b). We tried to persuade them this wasn't a very good idea (lower attitude scores in this experiment indicated more persuasion) with persuasive arguments allegedly provided by attractive or unattractive sources (using pictures from the Chicago Face Database).

Because evidence-based frames appeared to prompt central route processing in our first experiment, we expected Argument Quality to matter only when participants were presented with evidence-based frames of the arguments. But because Source Attractiveness typically operates as a peripheral route cue, we expected Source Attractiveness to matter only when participants were presented with bullshit frames of the arguments. As expected, Argument Quality made a difference—but only when arguments were provided in the evidence-based frames; it didn't matter when the arguments were provided in bullshit frames. What *did* matter in bullshit frames was the peripheral cue of attractiveness; attractive sources were more influential than unattractive sources when arguments were provided in bullshit frames, but the difference in attractiveness didn't appear to matter when arguments were provided in evidence-based frames.

Bullshit Sleeper Effect

We explored two additional ways in which bullshit might affect attitudes and beliefs about what is true. In our third experiment, we borrowed from a traditional sleeper effect procedure. A sleeper effect is a persuasive influence that increases, rather than decays, over time (Albarracín et al. 2017; Cook and Flay 1978; Priester et al. 1999). The easiest way to demonstrate a sleeper effect is by providing people positive information about a novel attitude object (e.g., political candidate) and showing that they express relatively positive attitudes (at least initially). Later, people are provided with contrary, negative information about the source of earlier information. In response, social perceivers tend to discount the initial information and adjust their attitudes downward (i.e., less positive). However, with nothing more than the passage of time, there tends to be an increase in persuasion such that attitudes return to a more positive valence. The sleeper effect is usually explained as a differential decay in memory—people forget the discounting cue sooner than they forget the initial persuasive message, such that the initial message retains a more persistent hold on attitudes (Pratkanis et al. 1988).

In our experiment, we presented over 200 participants with

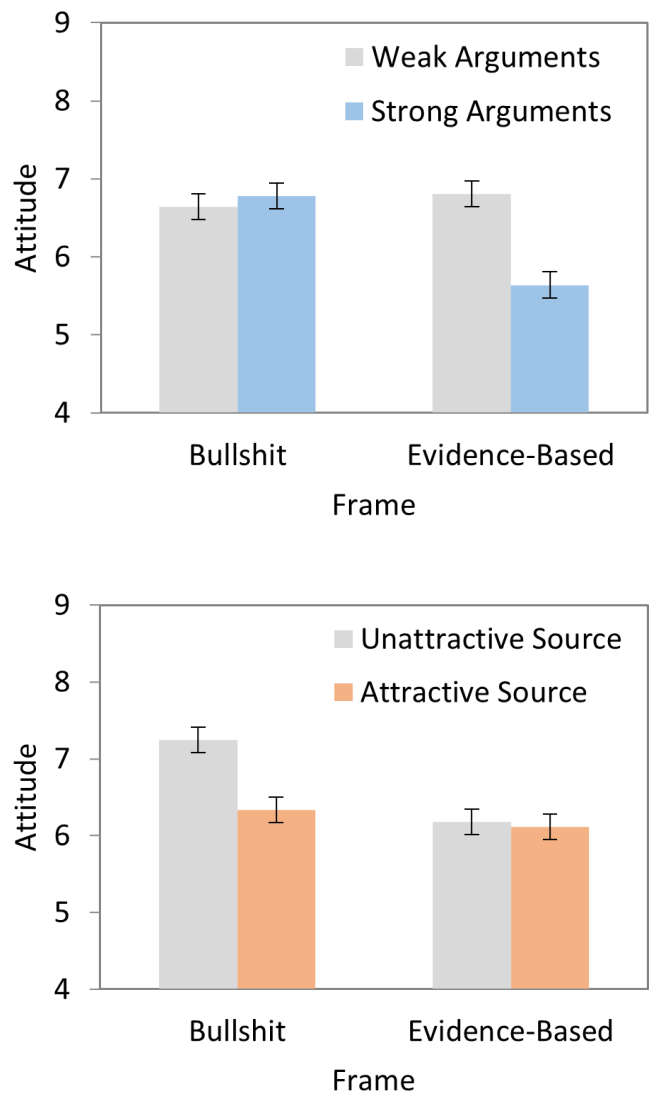


Figure 2. Thought Favorability means by Argument Frame and Argument Quality and by Argument Frame and Source Attractiveness; error bars represent ± 1 standard error. Lower scores indicate more persuasion that a longer Spring Break is a bad idea [Petrocelli 2021b].

an advertisement promoting multiple, desirable aspects of a fictitious pizza (Petrocelli, Seta, et al. 2023). Attitudes about the pizza were measured, and sure enough, they were overwhelmingly positive. Later on, participants were randomly assigned to one of three discounting cue conditions—one-third were informed that a consumer protection agency revealed that some of the information in the advertisement contained lies; one-third were informed that the advertisement contained bullshit; and one-third received no discounting cue at all. Attitudes about the pizza were measured once again, and, as expected, the attitudes of participants who received lies or bullshit were significantly reduced (see Figure 3). When our participants returned to our lab ten to fourteen days later, they were reminded of the attitude object we asked them to read about. Then we measured their attitude toward Ciao's Pizza one last time.

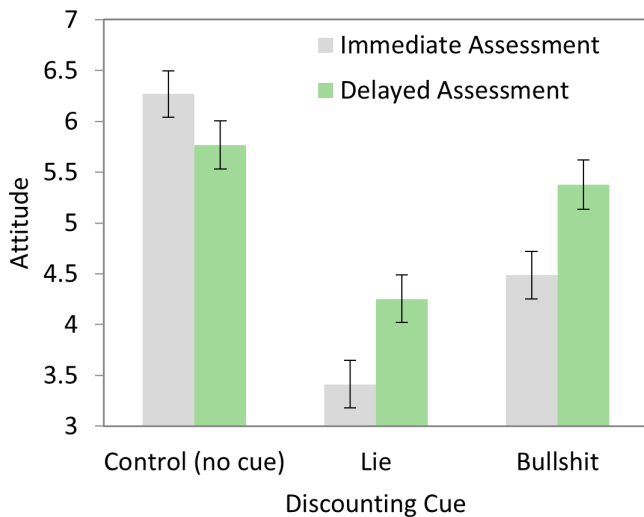


Figure 3. Attitude means for Attitude Assessment by Discounting Cue Condition [Petrocelli, Seta, et al. 2023].

After the delay in time, attitudes in both the lie and bullshit conditions increased. However, what we found most interesting was consistent with what we'd expect from the insidious bullshit hypothesis: Not only did the bullshit discounting cue result in more positive attitudes over time compared to the lie discounting cue, but it returned attitudes to a level of positivity that we would expect if our participants had never been given a discounting cue at all—and in a way that the lie discounting cue did not. Because bullshitting appears to produce a stronger sleeper effect for unwarranted attitudes than does lying, and attitudes are relevant to thinking and behaving, it serves as further support for the insidious bullshit hypothesis.

Bullshit Illusory Truth Effect

Our final experiment borrowed from a procedure used to demonstrate an *illusory truth effect*—the phenomenon that repeated statements are easier to process and remember (i.e., processing fluency) compared to new statements, regardless of whether repeated statements are truthful (Hasher et al. 1977). When something becomes *perceptually fluent*, it is more easily mentally processed. If we've seen something before, it's easier to recognize and easier to mentally process, and we often take the relative ease to signal truth.

It is relatively easy to produce an illusory truth effect (sometimes referred to as a repetition-induced memory) by first exposing people with everyday stimuli (e.g., names, statements) repeatedly. If an individual reads or hears the statement “Our sun is a red giant star” many times over, they are more likely to believe that statement than the truthful statement “Our sun is a yellow dwarf star.” The initial exposure sets contrasts with new information not previously processed; in other words, previously processed information is more perceptually fluent than new information. This “difference” in perceptual fluency is confused for truth. The illusory truth effect can be pushed even further, as empirical research has

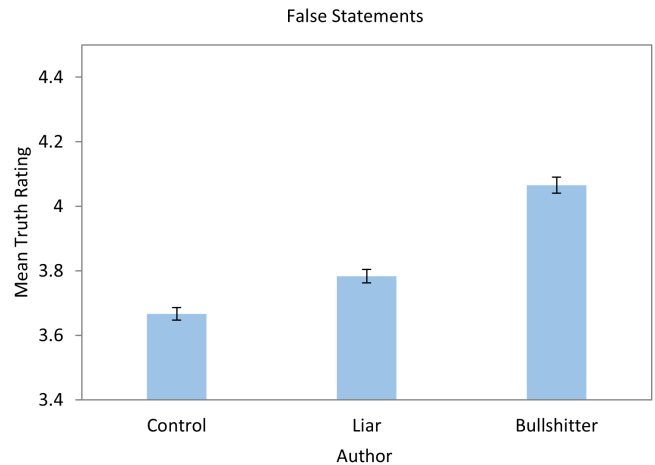


Figure 4. Mean truth ratings by Author Condition [Petrocelli, Rice, et al. 2020].

shown that even when individuals know better, they still fall victim to the phenomenon (Fazio et al. 2015).

In our final experiment, we used stimuli such as “Steinbeck is the last name of the author of *The Agony and the Ecstasy*” (false; it's Stone). We first asked over 300 participants to simply rate how interesting they found each statement for the purpose of exposing them to some information that would be repeated again later. Later, we asked participants how true they believed similar statements were. Some of the statements were repeated, some were new, some were false, and some were true. Our participants were randomly assigned to one of three conditions in which they learned 1) the author of the statements was instructed to construct statements whereby they knew half were true and half were false (liar); 2) the author was instructed to include statements whereby they knew half were true and half were without any real concern for truth, genuine evidence, or existing knowledge (they neither knew nor cared if they were true; a bullshitter); or 3) the author was simply instructed to compile a list of statements and half the statements were true and half were false (control).

When our participants were informed that some of the information was false after the initial exposure, we observed a drop in truth ratings for the control and liar conditions that was not observed for the bullshitter condition (see Figure 4). Once again, these results are consistent with the insidious bullshit hypothesis. Not only were all the items false, but bullshit exposure did not afford the same corrective attempt observed by the liar and control conditions. The very same pattern was observed for the items that happened to actually be true. There was an overall increase in truth rating for true information—and that was good—but the differences between those exposed to bullshit, versus honesty and lies, remained.

Despite the unfortunate actions and consequences of seemingly so many victims of bullshit (e.g., Edgar Welch, the Washington, D.C., Comet Ping Pong pizzeria gunman), people often think bullshit is harmless. Yet data from our Bullshit Studies Lab has demonstrated through four controlled experiments that bullshit can be more influential than lies



and—under some conditions—even evidence-based information. These results provide support for Frankfurt’s original insidious bullshit hypothesis, as well as a preliminary answer to the why question (i.e., why people engage in so much bullshitting behavior). Although bullshit may serve as a detriment to those who receive it, it can work in favor of the individual bullshitter. A just-desert would appear to be—if it be true—that bullshit may get one to the top, but it never lets one stay there (Sadhguru 2016). Downstream, negative consequences for bullshitters is a topic that awaits empirical attention. Until then, it appears that a little concern with the truth, access to readily available evidence, and perhaps treating bullshit as being as harmful as lies—until otherwise supported by evidence—will make all the difference. ■

Note

1. This article was adapted from a presentation at CSICon on October 22, 2022, in Las Vegas, Nevada.

References

- Albarracín, D., G.T. Kumkale, and P.P.-D. Vento. 2017. How people can become persuaded by weak messages presented by credible communicators: Not all sleeper effects are created equal. *Journal of Experimental Social Psychology* 68: 171–180.
- Cheyne, J.A., and G. Pennycook. 2016. The seductions of pretentious bullshit: An empirical study. *Skeptic* 21(1): 40–45.
- Cook, T.D., B.R. and Flay. 1978. The temporal persistence of experimentally induced attitude change: An evaluative review. In L. Berkowitz (ed.), *Advances in Experimental Social Psychology*, vol. 11. New York, NY: Academic Press.
- Fazio, L.K., N.M. Brashier, B.K. Payne, B.K., et al. 2015. Knowledge does not protect against illusory truth. *Journal of Experimental Psychology: General* 144: 993–1002.
- Frankfurt, H.G. 1986. On bullshit. *Raritan Quarterly Review* 6: 81–100.
- Hasher, L., D. Goldstein, and T. Toppino. 1977. Frequency and the confer-

- ence of referential validity. *Journal of Verbal Learning and Verbal Behavior* 16: 107–112.
- Littrell, S., E.F. Risko, and J.A. Fugelsang. 2020. The Bullshitting Frequency Scale: Development and psychometric properties. *British Journal of Social Psychology* 60: 248–270.
- Petrocelli, J.V. 2018. Antecedents of bullshitting. *Journal of Experimental Social Psychology* 76: 249–258.
- . 2021a. *The Life-Changing Science of Detecting Bullshit*. New York, NY: St. Martin’s Press.
- . 2021b. Bullshitting and persuasion: The persuasiveness of a disregard for the truth. *British Journal of Social Psychology* 60: 1464–1483.
- Petrocelli, J.V., E.N. Rice, and S.X. Shang. 2020. *Testing the Insidious Bullshit Hypothesis with the Illusory Truth Effect*. Poster session presented at the 21st annual convention of the Society for Personality and Social Psychology, New Orleans, LA.
- Petrocelli, J.V., C.E. Seta, and J.J. Seta. 2023. Lies and bullshit: The negative effects of misinformation grow stronger over time. *Applied Cognitive Psychology* 37: 409–418.
- Petrocelli, J.V., H.E. Silverman, and S.X. Shang. 2023. Social perception and influence of lies vs. bullshit: A test of the insidious bullshit hypothesis. *Current Psychology* 42: 9609–9617.
- Pratkanis, A.R., A.G. Greenwald, M.R. Leippe, et al. 1988. In search of reliable persuasion effects: III. The sleeper effect is dead: Long live the sleeper effect. *Journal of Personality and Social Psychology* 54: 203–218.
- Priester, J., D. Wegener, R. Petty, and L. Fabrigar. 1999. Examining the psychological process underlying the sleeper effect: The elaboration likelihood model explanation. *Media Psychology* 1: 27–48.
- Sadhguru. 2016. *Inner Engineering: A Yogi’s Guide to Joy*. New York, NY: Harmony.



John V. Petrocelli is professor of psychology in the Department of Psychology at Wake Forest University. He is an experimental social psychologist and is the author of *The Life-Changing Science of Detecting Bullshit* (St. Martin’s Press, 2021). He also serves as an associate editor for the *Journal of Personality and Social Psychology* and *Personality and Social Psychology Bulletin*.