

# Bullshit and its Detection

## *The Conditions Under Which Common, Everyday Bullshit is Most Likely to Emerge*

BY JOHN V. PETROCELLI

BULLSHITTING INVOLVES INTENTIONALLY OR unintentionally, consciously or unconsciously, communicating with little to no regard or concern for truth, genuine evidence, and/or established semantic, logical, systemic, or empirical knowledge.<sup>1,2</sup> As such, bullshitting is often characterized by, but not limited to, using rhetorical strategies designed to disregard truth, evidence and/or established knowledge, such as exaggerating or embellishing one's knowledge, competence, or skills in a particular area or talking about things one knows nothing about in order to impress, fit in with, influence, or persuade others.

Bullshitting is not simply the act of engaging in casual conversation, nor is it the same as lying. Lying involves actually knowing and caring about the truth. The bullshitter doesn't know what the truth is, doesn't care what the truth is, and isn't even trying to know. In fact, what the bullshitter says may actually be correct, but he wouldn't know it because he isn't paying any attention to truth or evidence for his claims. Bullshitting also has more to do with intentions than content. I may say "Pluto is a planet" and you may say "Pluto is a not planet." But if I don't care about the truth of that statement, and you on the other hand have considered the definitions of a planet, etc., then I'm probably bullshitting and you are not.

Like liars, bullshitters appear to be genuinely concerned with the truth and often believe their own bullshit, and these are things that make bullshit so difficult to

detect. To pull off a successful lie, the liar distorts his portrayal of reality and tries to remember his lie. The bullshitter doesn't have these cognitive burdens because most often he actually believes his own bullshit. Think of how much easier it would be to lie if we didn't have the burden of knowing the truth. It wouldn't feel much like lying at all. But the reason I'm more concerned about bullshitting behavior is that bullshit—because it has no connection to a concern for truth—can be more dangerous than lies.

We already know that psychic readers, alternative medicine enthusiasts, hypnotists, faith healers and the like, fabricate things in order to gain fame, money and power. But what are the sources of the more common sorts of everyday bullshit we continually face, like the communications we have with our families, friends, businesses, organizations, or any place in which people work together, share information, and make decisions? Under what condi-

tions will we be exposed to bullshit when we need information from people who are trying to sell us things, like a used car dealer, a wine seller, a jewelry dealer, or a realtor, or those who are not necessarily trying to directly sell us anything but may hope to influence our perspectives or decisions nonetheless, such as our friends and family, colleagues, doctors, and politicians?

Importantly, I'm not talking about senseless *old-school bullshit* as in corporate/business speak, for



example, “Our collective mission is to functionalize customer-driven enterprise solutions for leveraging underutilized portfolio transparencies and the bandwidth of our benchmark phlogistic sales.” Nor the content found in alternative medicine-quackery that makes senseless bullshit sound profound, such as “Hidden meaning transforms unparalleled abstract beauty.” What I am talking about is the more sophisticated *newfangled bullshit* more commonly encountered. Newfangled bullshit has meaning and prescriptive implications, incorporates numbers and figures, and is usually designed to inform, influence, persuade, or to convince people that we know what we are talking about when we really do not. Newfangled bullshit may be inaccurate, incorrect, or false, as in “Consumers should be aware of plastic rice from China that is difficult to discern from real rice as up to 15% of rice from China contains plastic.” Or it may be accurate, correct, or true, as in “Food stamp fraud is at an all-time high with \$70 million of taxpayer money was wasted in 2016 according to the USDA.”<sup>3</sup>

### The Problems with Bullshit

Bullshitting has important consequences that reach well beyond the risks of managing positive interpersonal impressions. Maybe we are not exactly living in a “post-truth world” in the literal sense that people are completely irrational beings with absolutely no attention to truth.<sup>4</sup> Yet, there remain rivers of bullshit for us to wade through on a daily basis, and it gets only in the way of optimal understanding and decision making.

A second societal problem stemming from the habit of bullshitting is the lack of bullshit detection skills. It is well known that if you have heard bullshit enough times, you begin to believe it. Many people continue to insist that giving children sugar makes them hyperactive, despite the evidence suggesting that sugar does not cause hyperactivity. And, how many of us still believe that Vitamin C is an effective treatment for a cold, despite the fact that science has demonstrated that there is no evidence of the link? Even the very statement “If you have heard bullshit enough times, you begin to believe it” is a myth. Research shows that you only need to hear it once!<sup>5</sup> Suppose someone told you that “Sydney is the capital city of Australia” or that “Styrofoam was invented in Norway” or that “Weapons of mass destruction are produced in Freedonia.” Perhaps you are wise enough to know that all of these statements are false, but you are more likely to misremember them as true simply because someone has suggested them to you. If we’ve

heard something before (even just once), our brains subconsciously use it as an indication that it’s probably true, and not an indication that it may be false.

Furthermore, cognitive psychologists who study attention, memory, and how we process information tell us that expertise and knowledge don’t always save us from the unwanted effects of bullshit. Consider the 4,800 investors who put good money into Bernie Madoff’s Ponzi scheme. His hundreds of investors included banks, investment firms, institutions, and pension funds, all with highly educated and sophisticated people managing them. If educated people continue to exchange genuine evidence and readily available information for bullshit, then we ignore our very best snapshots of reality. Without reality, we simply cannot make good decisions. Better information does not always lead to better decision making, but better decision making almost always requires better information. And, information full of bullshit does not make for an accurate picture of reality.

### Why Do People Bullshit?

One obvious reason that people bullshit is that they expect it to work. That is, bullshit is used to influence and get the bullshitter what he/she wants. Bullshitting is also more convenient than lying. Not only does the bullshitter not have the hassle of discovering the truth, but the social consequences of bullshitting are not as severe as those of lying.

Bullshit would be especially useful if it was more influential than lies. Because people tend to reject and dismiss liars faster and more reliably than they do bullshitters, one possibility is that underestimations of the insidiousness of bullshit ironically places bullshit in a more potent position to influence than lies. My *Bullshit Sciences Lab* at Wake Forest University where I am a professor in the psychology department conducted an experiment to determine if this is the case. In one experiment, we utilized a classic sleeper effect procedural paradigm. The *sleeper effect* is anything that stands as a persuasive influence that increases, rather than decays, over time.<sup>6</sup> The easiest way to demonstrate a sleeper effect is to have individuals mentally process positive information about a novel attitude object (e.g., a new brand of frozen pizza) in the form of a persuasive message. Later, if people learn the source of the persuasive message was dishonest or incompetent, they discount the initial information and adjust their attitudes downward. However, with the passing of time (e.g., two weeks), people appear to experience a differential decay in memory for the initial positive information and the discounting

cue, such that mental representations for the initial persuasive arguments decay more slowly than those of the discounting cue (which occurs more rapidly), resulting in more positive attitudes among people who were lied to.

In our experiment, participants learned about a fictitious gluten-free pizza and were informed of its positive qualities.<sup>7</sup> Next, participants were provided with a discounting cue explaining that either parts of the advertisement contained lies or bullshit. Attitudes toward the pizza were then measured immediately afterwards as well as 14 days later.

Consistent with previous sleeper effect studies, attitudes generally became more positive about the pizza from the time of the immediate assessment to the delayed assessment. However, our bullshit condition participants reported significantly more positive attitudes than their lie condition counterparts immediately as well as after a 14-day delay. In other words, the sleeper effect was significantly stronger for bullshit than that which we and others have demonstrated for lies.

### **Measuring Bullshitting and Bullshit Detection**

People find bullshitting to be an incredibly easy task. In fact, individuals will bullshit about anything—a convenient reality when attempting to study bullshitting behavior.

To measure bullshitting behavior, all that is necessary is to first ask individuals to report their attitude about something (e.g., nuclear power) and to then provide reasons for their attitudes. Once they provide reasons for their attitudes, they are asked to rate each thought listed with regard to how much they were truly concerned with genuine evidence or existing knowledge when they listed the thought.

Measurement of bullshit detection is about as simple as measuring bullshitting. Gordon Pennycook and his colleagues have provided a useful method.<sup>8</sup> All that is necessary is to present individuals with a number of meaningless, but syntactically correct, statements containing impressive sounding scientific, philosophical, and metaphysical jargon (e.g., “We are in the midst of a self-aware blossoming of being that will align us with the nexus itself.”) and ask them to provide profundity ratings. Profundity ratings of the senseless statements can be juxtaposed with profundity ratings of prototypically profound statements (e.g., “A river cuts through a rock, not because of its power but its persistence.”). Ratings of both sets of statements can be combined to form general *receptivity*

to bullshit and sensitivity to bullshit measures.

### **The Psychology of Bullshitting**

What makes someone particularly likely to bullshit? This is an important empirical question because understanding the conditions under which bullshit is most likely to emerge is fundamental to calibrating our bullshit detectors and deflecting its unwanted effects.

Harry Frankfurt, the philosopher who articulated the definition of bullshit, formulated two important causes of bullshitting that have been studied in my Bullshit Sciences Lab. First, Frankfurt asserted that “Bullshit is unavoidable whenever circumstances require someone to talk without knowing what he is talking about. Thus the production of bullshit is stimulated whenever a person’s obligations or opportunities to speak about some topic are more extensive than his knowledge of the facts” (p. 99). Indeed, bullshitting appears to be predicated on the assumption that people feel an obligation to have/provide an opinion about everything (from apples to zebras); but of course, people cannot have an informed opinion about everything. People use bullshit to express their opinions (informed or ill-informed), provide others with information, or “try out” what it feels like to express such opinions when it is clear that the social expectations to have an opinion are relatively great.

I tested these ideas in a simple experiment in which participants were told that the experiment was designed to better understand how people explain the behaviors of others. They were told that there is a person named Jim. Half of the participants learned some additional things about Jim’s personality whereas the other half did not. It was explained that Jim was running for public office and made some decisions, and that we were interested in their explanations for his decisions. Respondents could list any thoughts they generated. Here is the critical piece—participants assigned to a condition obligated to provide an opinion were given no additional instructions for the thought-listing task, whereas participants assigned to an unobligated condition were extensively reminded that they were under no obligation to provide an opinion and could skip the thought-listing task altogether if they elected to.

The experiment resulted in four different groups of respondents, those who were either knowledgeable or unknowledgeable of Jim and either obligated or not obligated to share their opinions of his decisions.

Interestingly, all participants listed at least a couple of thoughts explaining Jim's decisions. This was particularly surprising because most survey respondents simply want to finish the survey and move on with their lives—and remember, half of the participants were not under any obligation to complete the most time-consuming part of the survey. Bullshit was 24% among unobligated respondents but almost twice that among obligated respondents (44%). It's disturbing that a third of our judgments (or that which we are willing to share) may be based on bullshit. Furthermore, the bullshit of participants unknowledgeable of Jim was 36% whereas it was only slightly smaller (33%) among participants knowledgeable of Jim. More surprising was that the knowledgeable/unknowledgeable percentages did not depend at all on the obligation to provide an opinion. That is, people who have some or no knowledge about a topic appear just as willing to bullshit when they are under no obligation to provide their opinions as when they are. People often feel obligated to talk about things they really know nothing about—and what often comes out is bullshit.

Frankfurt further theorized that no matter how “studiously and conscientiously the bullshitter proceeds, it remains true that he is also trying to get away with something” (p. 87). As long as they are not being lied to, people appear to be tolerant of bullshit. It follows that people engage in bullshitting when they anticipate ease in receiving a “social pass” of acceptance or tolerance for communications characterized by little to no regard for how things really are.

One context in which it is easy to get away with bullshit is one in which relatively few people appear to have an informed opinion. If most people do not appear to have an informed opinion (i.e., social expectations to have an informed opinion are weak), getting away with bullshit should be easy.

In my experiment involving explanations of Jim's behavior, I included additional conditions that enabled a closer look at the role of the ease of getting away with bullshit in encouraging the production of bullshit. Participants were also either led to believe that their thought-listings would be evaluated for their accuracy by other people who knew Jim very well or by people who did not know Jim at all.

Confirming the idea that people are relatively more likely to bullshit when they expect to get away with it, bullshit was significantly greater (41%) when participants were led to believe their thoughts would be evaluated by unknowledgeable coders than when they were led to believe their thoughts would be evaluated by knowledgeable coders (29%). Interestingly, the

link between bullshitting and the relative ease of getting away with it depended in part on the obligation to provide an opinion. Specifically, bullshitting was higher (37%-45%) compared to when participants were knowledgeable and did not have any obligation to share their opinions (15%). Bullshitting is relatively prevalent when it is clear that the audience signals either a lack of concern for evidence or no demand for it. That is, anything that signals to people that expectations to express evidence-based thoughts are low, or they are not required to know what they are taking about, will be followed by bullshit.

My Bullshit Studies Lab conducted two additional experiments suggesting another important antecedent of both bullshitting behavior and bullshit detection.<sup>9</sup> *Self-regulatory resources* involve the psychological assets people use to manage and control their own thoughts, feelings, and behaviors. Under conditions of high social accountability, people are relatively motivated to attenuate their contributions to bullshit. Yet, one's ability to attenuate his/her own bullshit should be negatively affected when self-regulatory resources are relatively unavailable, thereby contributing more bullshit relative to when self-regulatory resources are more available. Likewise, recognition, detection, and motivation to protect against the potential contamination of bullshit is a task requiring self-regulatory resources; thereby people should be less successful at detecting bullshit when their self-regulatory resources are relatively unavailable.

Within the dual-process model tradition associated with Amos Tversky and Daniel Kahneman, one of the hallmark distinctions between System 1 (intuitive) and System 2 (deliberative) processing is the reliance on self-regulatory resources: System 1 processing runs autonomously and does not require the expenditure of working memory capacity and self-regulatory resources, whereas System 2 processing is deliberate, effortful, and dependent upon working memory capacity and self-regulatory resources to be successfully executed.<sup>10</sup> Indeed, one of the most common ways by which researchers are able to distinguish whether a given outcome is due to System 1 or System 2 processing is to manipulate the availability of self-regulatory resources through a cognitive load manipulation. If the outcome is influenced by resource availability, the results indicate that System 2 processing is responsible for that outcome. On the other hand, if the outcome is unaffected by resource availability, results indicate that System 1 processing is responsible. Successful bullshit detection, as well as evidence-based communication in place of bullshit,

require the self-regulatory resources necessary to engage in System 2 processing.

To compare the bullshitting behavior and bullshit detection abilities of individuals depleted of their regulatory resources with individuals not depleted, one experiment first determined the chronically preferred time to perform cognitively effortful tasks of all individuals—some people prefer to work in the morning whereas others prefer to work in the evening—and then required them to perform experimental tasks during their preferred (not depleted) or non-preferred times (depleted). In another experiment, individuals were asked to write three short essays, for five minutes each. The catch was that individuals assigned to the *easy* writing task were instructed to not use the letters “x” and “z” anywhere in their essays (not depleted), whereas individuals assigned to the *difficult* writing task were instructed to not use the letters “a” and “n” anywhere in their essays (depleted) before measuring their bullshitting and bullshit detection abilities.

Both experiments revealed the same results. When individuals were relatively depleted of their regulatory resources, either by engaging in experimental tasks after becoming fatigued from prior tasks or being asked to complete a task during their chronically non-preferred time, they engaged in relatively more bullshitting and found it more challenging to detect bullshit. Evidence-based communication (i.e., no bullshit) and successful bullshit detection appear to require deliberate and effortful thought that harnesses available working memory capacity and self-regulatory resources.

### Conclusions

Progress in the struggle against bullshit is advanced by a deeper understanding of the conditions under which bullshit emerges. We now know that people engage in considerable bullshitting when social cues make them feel obligated to provide an opinion about something, even when it is something they know nothing about. We also know that people bullshit when they expect it to be relatively easy to get away with it. That is, when people detect cues from the social context suggesting it will be easy to get a “social pass” of acceptance or tolerance for their bullshit (like when we make it easy by failing to put up any barriers of intolerance). We also know that people are more likely to engage in bullshitting behavior, and are relatively less likely to detect bullshit, when they are relatively cognitively fatigued.

What can we do about the mountains of bullshit we face everyday, especially when some liveli-

hoods depend on bullshit (e.g., used car dealers)? What might life be like, and how might our world run, if rather than being the product of bullshit, judgments and decisions were based on evidence and truth? How can we change our communicative culture to reduce bullshit and its unwanted effects?

First, consider something you may have not considered before. You’re probably not going to believe this, but you are just as susceptible to bullshit as anyone else. It’s critical that we all entertain this possibility. In fact, a major source of the unwanted effects of bullshit is that everyone thinks they can readily detect bullshit and thereby feels unaffected by it, despite research clearly demonstrating that bullshit is not easily detected. Bullshit is not only the stuff of psychics who speak with the dead, Ouija boards enthusiasts, alternative medicine quacks, and hypnotists who retrieve repressed memories. Rather, bullshit can be found everywhere.

Second, we need to face the reality that part of the problem is that we often prefer bullshit over the truth. Believing bullshit can be comfortable. In fact, people have multiple motivations for believing bullshit—and part of that problem has to do with how we are wired. The confirmation bias is extremely strong. That is, people are strongly wired to attend to information that confirms their preconceived notions and ignore important disconfirming evidence. Likewise, our bullshit filters and detectors are not particularly well-tuned, especially when bullshit aligns with our views of the world. We want to believe we know things we really don’t. We want to believe we’ve made the right decisions: the car we bought, the career we pursued, the person we married, the school we sent our children to, the candidate we voted for. There is a feeling of security in thinking we are right. But, there is a liability in pretending we know things we really do not, and wanting things to be true doesn’t make them so.

Third, recall that people bullshit when they expect they will not be called to account for it. What would happen if we started calling bullshit more regularly and stopped making it so easy to get away with? Susceptibility to bullshit is best explained by our failure to ask important questions about the content we encounter. When people claim that a moon landing was faked, that a used car will likely go another 100,000 miles, that a politician can solve all of our problems, or when a Pollyannaish TED talk speaker tries to save the world again in only 18 minutes, we should ask ourselves: What agenda is behind this message? Who is providing the evidence? What is their credibility? And, we should ask the



potential bullshitter: Why do you think this? How do you know this to be true? And what sort of evidence supports your conclusion?

Of course, calling bullshit can be a real conversation killer—and we don't want to cut all lines of communication. Thus, a few rules should be in order for dealing with bullshit.

1. Don't call bullshit unless you are sure it is bullshit.
2. Consider the possibility that you are the one who is confused. Most people have found that even acting a bit confused can prompt the bullshitter to clarify his claim, and clarification is a major antidote to bullshit.
3. Attack the claim, and not the person.
4. Reduce confusion to an understandable error in reasoning; doing so is more forgiving and a gentler pill for the bullshitter to swallow.
5. When you find yourself guilty of bullshitting, just

admit the fault. We all make mistakes, so don't double down on your bullshit. It only makes things worse.

Even when conformity is its strongest, often times *it only takes one person* to stop the unwanted effects of bullshit. We know from research that the influence of what people think they should do can outweigh the power of social norms.<sup>11</sup> Armed with the power of evidence, and perhaps a few allies interested in the same, everyone can play an important role in the struggle against bullshit. We can blame Facebook, Twitter and web profiteers all we want for the escalation in bullshit, but at some level the responsibility comes down to each of us having the ability to search for truth, discern fact from fiction, and communicate what we know to be true, not just what we want or hope to be true. Competent bullshit detection and disposal depend on it. **S**

---

## REFERENCES

1. Frankfurt, H. 1986. "On Bullshit." *Raritan Quarterly Review*, 6, 81-100.
2. Petrocelli, J. V. 2018. "Antecedents of Bullshitting." *Journal of Experimental Social Psychology*, 76, 249-258.
3. Although true, the applied implications of this information is bullshit until one applies a bit of logic to see that \$70 million is less than one fifth of 1% of the entire cost of the food stamp program (Supplemental Nutrition Assistance Program); an incredibly efficient and laudable program.
4. Pinker, S. 2019. "Why We Are Not Living in a Post Truth Era: An (Unnecessary) Defense of Reason and a (Necessary) Defense of Universities' Role in Advancing it." *Skeptical*, 24(3), 26-33.
5. Fazio, L. K., Brashier, N. M., Payne, B. K., and Marsh, E. J. 2015. "Knowledge Does Not Protect Against Illusory Truth." *Journal of Experimental Psychology: General*, 144, 993-1002.
6. Hovland, C. I., Lumsdaine, A. A., and Sheffield, F. D. 1949. *Experiments on Mass Communication*. Princeton, NJ: Princeton University Press.
6. Hovland, C. I., and Weiss, W. 1951. "The Influence of Source Credibility on Communication Effectiveness." *Public Opinion Quarterly*, 15, 635-650.
7. Petrocelli, J. V., Seta, C. E., Seta, J. J., and Shang, S. X. Manuscript under review. "When Bullshitters are More Persuasive than Liars: Testing the Insidious Bullshit Hypothesis with the Sleeper Effect."
8. Pennycook, G., Cheyne, J. A., Barr, N., Koehler, D. J., and Fugelsang, J. A. 2015. "On the Reception and Detection of Pseudo-Profound Bullshit." *Judgment and Decision Making*, 10, 549-563.
9. Petrocelli, J. V., Watson, H. F., and Hirt, E. R. in press. "Self-Regulatory Aspects of Bullshitting and Bullshit Detection." *Social Psychology*.
10. Barbey, A. K., and Sloman, S. A. 2007. "Base-Rate Respect: From Ecological Validity to Dual Processes." *Behavioral and Brain Sciences*, 30, 241-297. Chaiken, S., and Trope, Y. (Eds.). 1999. *Dual-Process Theories in Social Psychology*. New York: Guilford Press. De Neys, W. (Ed.). 1999. *Dual Process Theory 2.0*. New York: Routledge. Evans, J., and Stanovich, K. E. 2013. "Dual-Process Models of Higher Cognition: Advancing the Debate." *Perspectives in Psychological Science*, 8, 223-241. Kahneman, D. 2011. *Thinking, Fast and Slow*. New York, NY: Farrar, Straus and Giroux. Sherman, J. W., Gawronski, B., and Trope, Y. (Eds.). 2014. *Dual-Process Theories of the Social Mind*. New York: Guilford Press. Shiffrin, R. M., and Schneider, W. 1977. "Controlled and Automatic Human Information Processing II: Perceptual Learning, Automatic Attending and a General Theory." *Psychological Review*, 84, 127-189. Sloman, S. A. 1996. "The Empirical Case for Two Systems of Reasoning." *Psychological Bulletin*, 119, 3-22.
11. Cialdini, R. B. 2006. *Influence: The Psychology of Persuasion*. New York: Harper Business; Cialdini, R. B., Reno, R. R., and Kallgren, C. A. 1990. "A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places." *Journal of Personality and Social Psychology*, 58, 1015-1026; Goldstein, N. J., Cialdini, R. B., and Griskevicius, V. 2008. "A Room With a Viewpoint: Using Social Norms to Motivate Environmental Conservation in Hotels." *Journal of Consumer Research*, 35, 472-482; Jacobson, R. P., Mortensen, C. R., and Cialdini, R. B. 2011. "Bodies Obligated and Unbound: Differentiated Response Tendencies for Injunctive and Descriptive Social Norms." *Journal of Personality and Social Psychology*, 100, 433-448.