

## WHAT'S WRONG WITH BULLSHIT

FLORIAN COVA

*Philosophy Department and Swiss Centre for Affective Sciences, University of Geneva*

Past philosophical analyses of bullshit have generally presented bullshit as a formidable threat to truth. However, most of these analyses also reduce bullshit to a mere symptom of a greater evil (e.g. indifference towards truth). In this paper, I introduce a new account of bullshit which, I argue, is more suited to understand the threat posed by bullshit. I begin by introducing a few examples of “truth-tracking bullshit”, before arguing that these examples cannot be accommodated by past, process-based accounts of bullshit. I then introduce my new, output-based account of bullshit, according to which a claim is bullshit when it is presented as or appears as interesting at first sight but is revealed not to be that interesting under closer scrutiny. I present several arguments in favor of this account, then argue that it is more promising than past accounts when it comes to explaining how bullshit spreads and why it is a serious threat to truth.

### 1. Introduction: Some Garden Varieties of Truth-Tracking Bullshit

Sebastian Marx is an American stand-up comedian who worked and lived in France, and loves to make fun of his experience there. In one of his shows, he recalls having worked in a French restaurant as a waiter, and mocks the “super complicated” names French restaurants sometimes give to their dishes. As an example, he mentions the “tasty fisherman’s hodgepodge on a bed of finely crushed puree”, before revealing: “it’s just fish and potatoes”. He then jokes about the fact that French people would say that such names are not “complicated” but “sophisticated”, and that Americans have another word for “sophisticated”: it’s... *bullshit*.

Sebastian Marx is far from being the first to make fun of such pompous dish names. French comic artist Gotlib has a cartoon in which he depicts someone at a restaurant ordering “the Prince of the Seas in his Sauce of the Fruits of Provence”

---

**Contact:** Florian Cova <florian.cova@gmail.com>

and looking disappointed when the waiter brings him sardines in oil. This is another paradigmatic case of bullshit. Still, it's an interesting one because it is true – and not even accidentally true: it actually *tracks truth*. Indeed, there is some connection between the name of the dish and its content, to the extent that trained consumers can guess what the dish contains. Rather than inventing names at random, the person who wrote the menu intended for the names to track the content of the dishes: they would never have given the name of “tasty fisherman’s hodgepodge on a bed of finely crushed puree” to a piece of deer with green vegetables (this would rather be “the hunter’s delight on its farandole of greeneries”).

Such mundane cases of “truth-tracking” bullshit are not limited to pompous restaurant dishes. We also use it to paint ourselves in a flattering light on our resume: we do not want to be caught telling something false, but we want our past professional experience to read as much more impressive than it actually was. You spent two months entering the results of pen-and-paper surveys into an Excel sheet in exchange for credits? Well, just put on your resume that you were “Responsible for data management, integration and integrity at Pr. Academic’s laboratory” (and do not forget to add “Advanced proficiency in Microsoft Excel” in the “Skills” section). You once gave a public speech in the street just in front of the UNO building? Just say that you presented your work in front of the UNO (Aberkane 2021). Again, these are not technically *false* claims – rather, they are fancy ways of presenting the truth.

In the same way that fancy dish names are clear cases of *bullshit*, fancy resumes are a paradigmatic case of *bullshitting*. And both cases seem to point at one crucial feature of bullshit: bullshit is something that sounds impressive at first sight but can be easily “deflated” on closer inspection. However, I will argue that existing philosophical accounts of bullshit have largely ignored these paradigmatic examples of bullshit.<sup>1</sup> In this paper, I will present a new account of bullshit that takes these examples into account and argue that this allows us to better understand what exactly the problem is with bullshit.

## 2. Limitations of Past Accounts of Bullshit

The examples I put forward have two purposes. The first is to highlight some limitations of former philosophical accounts of bullshit. As I will argue in this section, the most prominent philosophical accounts of bullshit fail to account for paradigmatic cases of “truth-tracking” bullshit.

---

1. West & Bergstrom (2020) offer a lot of additional examples of truth-tracking bullshit, such as graphs that are literally true but present data in a deceptive manner.

### 2.1. Frankfurt's account

In his essay *On Bullshit*, Frankfurt (2005) characterizes *bullshit* as “an indifference to truth”. More precisely, Frankfurt argues that one is bullshitting when (i) one does not care about the truth of what one is asserting, and (ii) one is trying to deceive us by hiding this fact from us.

A good illustration of Frankfurt's theory would be the case of French psychoanalyst Jacques Lacan. In 1973, Lacan closed a TV interview with these two enigmatic sentences: “*L'interprétation doit être prête pour satisfaire à l'entreprêt. De ce qui perdure de perte pure à ce qui ne parie que du père au pire*” (“The interpretation must be ready to satisfy the undertaking. From that which lasts of pure loss to that which only bets on the father at worst”). According to anecdotal reports, when later asked by a colleague what he meant by that, Lacan simply replied “*J'ai dit ça pour les assonances*” (“I said that for the assonances”) (Van Rillaer 2019). Supposing this is true, then Lacan's statements are a perfect instance of bullshit as defined by Frankfurt: Lacan did not care what these statements meant, even less whether they were true or false – only whether they sounded good. Hence, he was bullshitting.

However, Frankfurt's account cannot accommodate the examples of “truth-tracking” bullshit I have presented. We saw that fancy dish names are clear, paradigmatic cases of bullshit: but can we say that the person who was hired to come up with these names and write the menu was indifferent to the truth of what they were writing? Surely not. They carefully crafted the names so that they would correspond in some ways to the content of the dishes. The same is true for bullshit CVs: those are crafted with an eye on truth, in order to escape charges of lying. So, we can say that, in both examples, the person who produced the bullshit was very concerned with what was true or not. Still, these are clear cases of bullshit.<sup>2</sup>

Of course, one might argue that those people do not care about truth *per se*, but only for ulterior motives: they care about the truth of their statements because they fear the reactions of angry customers or accusations of lying. One might thus be tempted to amend Frankfurt's account so that one is said to be bullshitting when one does not care about the truth of one's statements *per se*. However, this seems to be too broad a definition. Imagine a student at an exam who gives correct answers, but cares more about having a good grade than about the truth

---

2. One might argue that Frankfurt's account can be slightly modified to accommodate these cases, by pointing out that bullshitters might not be indifferent to the truth of the proposition they explicitly express, but might still be indifferent to the truth of the proposition they try to signal or communicate implicitly. However, I don't think such an amendment would work: even if the person composing the menu was intentionally trying to mislead people by giving an overly flattering image of the dishes, the names would still be *bullshit*. (Moreover, in the menu case, it is not clear that any precise proposition is implied, besides a general impression. I will get back to this point in §5.)

of their answer (or a student who is a Young Earth creationist but still gives the answer their biology teacher expects without believing them to be true): it would be strange to say that this student is bullshitting and that his answers are bullshit. We can also imagine a scientist who desperately wants to win a Nobel Prize, and for which securing the associated recognition is more important than truth itself. However, he knows that his theories and discoveries need to be true for him to win the Prize. So, he works hard, scrupulously scrutinizes the evidence he collects, and amends his theories so that they track the truth. Is this scientist *bullshitting* whenever he presents his theories to his colleagues? And can we really say that his theory is *bullshit*? This sounds very strange.

Moreover, going this way will lead us to consider that most (if not all) cases of lies are cases of bullshit: indeed, the person who lies can hardly be considered to care about the truth of their statements *per se*, as we often lie for ulterior motives. Thus, the only case in which a person who lies would not also be bullshitting would be “Augustinian” cases in which the agent lies for the pleasure of lying – which, admittedly, cannot be considered the most paradigmatic cases of lying. But Frankfurt stresses the fact that lying is not the same as bullshitting. Thus, it seems that demanding that the agent cares about the truth *per se* in order for them to not be bullshitting is not a satisfying solution.

It is important to note that I am not the first to present cases of bullshit produced by agents who are not indifferent to the truth of their statements. Carson (2016) already presented several counter-examples along these lines. One of his examples is the case of a politician who was asked whether she would be willing to nominate people to the US Supreme Court who support *Roe vs. Wade*. Because she is anti-abortion but does not want to lose potential voters who support the right to abortion, the politician gave an answer that fails to answer the question she was asked:

There are many things that need to be taken into account when nominating someone for the Supreme Court. This isn't the only relevant consideration. Many other factors are also important. The Supreme Court is a venerable institution and, as our Founding Fathers wisely intended, a central pillar of our blessed democracy. We need outstanding people to sit on the Supreme Court. I would nominate someone with an outstanding intellect and legal mind who has adequate judicial experience and supports my judicial philosophy of following the Constitution as it is written. (Carson 2016: 57)

This is a case of “evasive bullshit” – the kind of bullshit we produce when we are pressured to answer questions that we do not want to answer. What is important is that, in such cases, the agent cares about saying something true: for example,

the anti-abortion politician does not want to be caught saying something false on TV. Still, it seems that she is *bullshitting*, as their answer is completely uninformative relative to the questions she was asked. As such, these cases provide additional reasons to reject Frankfurt's account.

The key difference between my examples and Carson's is that my examples are not cases of "evasive bullshit": they provide new information that is relevant to the question under discussion ("what's on the menu?", "what is your professional experience?"). This will prove important as we move on to accounts of bullshit and *bullshitting* that try to improve on Frankfurt's proposal in order to better accommodate Carson's examples.<sup>3</sup>

## 2.2. Stokke and Fallis' Account

Taking note of Carson's counter-examples, Stokke and Fallis (2017) argue that Frankfurt's approach should be integrated into a broader account according to which "bullshitting is a mode of speech marked by indifference toward inquiry, the cooperative project" (277). More precisely, they argue that "the bullshitter is characterized by a specific kind of indifference toward inquiry," where inquiry in discourse is defined as the "cooperative project of incremental accumulation of true information with the aim of discovering how things are, or what the actual world is like" (279). Drawing on work by Craige Roberts (2012), Stokke and Fallis consider that each instance of inquiry is characterized and structured by a set of "questions under discussion," e.g. "when does the bus leave?," "how cold is it outside?," "has John done his homework?" They then define *bullshitting* as "a mode of speech marked by indifference toward contributing true or false answers to questions under discussion" (2017: 279). Thus, the kind of indifference toward truth or falsity that characterizes the phenomenon of *bullshitting* is not indifference toward the truth-value of what one says, but indifference toward the effect that one's contributions have on the discourse.

Drawing on this framework, they put forward the following definition of *bullshitting*:

*A* is *bullshitting* relative to a [question under discussion] *q* if and only if *A* contributes *p* as an answer to *q* and *A* is not concerned that *p* be an answer to *q* that her evidence suggests is true or that *p* be an answer to *q* that her evidence suggests is false. (2017: 295)

---

3. Meibauer (2016) draws on pragmatic theory to improve Frankfurt's account. However, these modifications do not allow Frankfurt's account to accommodate Carson's (2016) counter-examples. Meibauer rather argues that these are not cases of bullshit, since they are cases in which the agent is not indifferent to truth.

According to Stokke and Fallis, this account of bullshitting allows one to accommodate Carson's examples of "evasive bullshit." Indeed, people who have recourse to evasive bullshit typically do so to avoid answering the question – and so, Carson's bullshitters are not concerned that their assertion be an answer to the question under discussion, which makes their answers instances of bullshitting.

But what about the examples of "truth-tracking bullshit" I put forward in the introduction? In such cases, it seems that the agents are concerned that their assertion be an answer to the question under discussion. Indeed, menus answer the question they are supposed to answer ("what's on the menu?"), and so do CVs ("what's the candidate's past experience?"). Thus, they are concerned that their assertion be an answer to the question under discussion. But are they concerned about what their evidence suggests about the truth or falsity of their statements?

As stated earlier, both the person who came up with fancy dish names for the menu and the person who composed a flattering resume were very concerned not to state anything that could be considered blatantly false. In fact, they were very concerned that their statements be true, *at least in one sense*. Thus, Stokke and Fallis' account of bullshitting does not properly capture what makes these examples instances of bullshitting, mainly because it is tailored to accommodate evasive bullshit, which they are not.

Not only does Stokke and Fallis' account produce false negatives by failing to detect such obvious instances of bullshitting as bullshitting, I also think it produces false positives by flagging as bullshitting cases that should not be considered as such. Take the following example: a student is taking an exam in philosophy on some Great Philosopher. He has heard that his teacher's interpretation of this Great Philosopher's claim is not shared by other specialists and commentators. Still, he does not care. He makes no effort to investigate other interpretations and does not really care whether his teacher's interpretation is true: he only cares about succeeding at the exam and never hearing about this Great Philosopher again. When the exam comes, he answers what he knows the teacher wants him to answer and succeed in the exam. Did he bullshit his way through the exam? Stokke and Fallis' account would predict that he did, but this sounds very strange.<sup>4</sup>

Take another example: Bob is suffering from a medical condition that makes him alternate between very good and very bad days. Because of that, people around him worry about him and always ask him how well he is. Bob has grown tired of seriously answering these questions, so he has taken the habit of always

---

4. Stokke and Fallis could avoid such difficulties by claiming that answering a question on an exam does not really constitute a case of participating in an inquiry by answering a question under discussion. However, doing that would make it impossible for them to accommodate Carson's student case (see Carson 2016: 59).



answering “fine” without giving much thought about how well he is. According to Stokke and Fallis’ account, Bob should be considered to be bullshitting. But, here again, that does not seem like a good example of bullshitting. (In fact, Stokke and Fallis’ account would predict that we are bullshitting every time we automatically answer “fine” to someone who asks us how we are faring without giving it much thought.)

### 2.3. From Johnson's Account to Moberger's Accounts

But Stokke and Fallis (2017) were not the only ones to attempt to improve Frankfurt’s account. Others have tried to modify Frankfurt’s account to emphasize epistemic elements rather than pragmatic ones. One motivation for these accounts is that, according to their promoters, it is possible to produce bullshit while being genuinely motivated to find and tell the truth. Most of their examples draw on people who endorse pseudoscientific theses without being charlatans: the idea is that at least some partisans of astrology (Moberger 2020), intelligent design (Reisch 2006), or climate change skepticism (Johnson 2010) genuinely believe in the claims they are making. Nevertheless, those claims are still bullshit. Johnson (2010) calls such instances of bullshit “culpably confused bullshit” and claims that they constitute one more counter-example to Frankfurt’s account—a claim with which I agree.

To account for such cases, these researchers have put forward accounts of bullshit that emphasize epistemic shortcomings in the attitudes of those who produce it. Johnson thus argues that “the essence of bullshit is not, as Frankfurt would have it, indifference to truth but a blamably insufficient deference to truth, a kind of intellectual sloppiness, a ‘culpable intellectual negligence’” (2010: 28). More precisely, he defines bullshit in the following way (2010: 14):

A person is guilty of bullshitting when, and only when:

1. the person implicitly or explicitly asserts a proposition;
2. in the assertion of the proposition, or in the prior process of coming to accept the proposition, any impartial interest in what is true is subordinated to or supplanted by a competing interest;
3. there is no exculpatory reason for such subordination or supplantation; and
4. if the proposition is explicitly asserted, it is not believed to be false.

However, it is not clear that this account can accommodate my cases of “truth-tracking” bullshit, as the bullshitters in these cases display a certain interest in truth. Of course, this interest is instrumental, but Johnson explicitly states that

an instrumental interest in truth is enough for avoiding bullshit (as restricting the account to non-instrumental interest for truth would have his theory run into the same difficulties I mentioned in section 2.1, and would also lead Johnson to conclude that the student who just cares about passing his exam would be bullshitting).

Moberger (2020) distinguishes between *indifference* towards the truth of one's statements, and *lack of conscientiousness* with respect to the truth of one's statements. The difference between the two is that one might care about one's statements being true (and thus not be indifferent towards truth) while failing to take the necessary steps and precaution to ensure that one's statements are true. Thus, people who believe in astrology without being profit-seeking charlatans *do* care about whether their statements are true (and are thus not indifferent to the truth of their statements), but they do not engage in the kind of inquiry that would be required to ensure that their statements are true. Moberger thus argues that *lack of conscientiousness* rather than *indifference* is the condition for bullshit.

How does Moberger's account fare when confronted with the examples of "truth-tracking bullshit" I presented in the introduction? Not that well, it seems to me. After all, it seems that the person who wrote the fancy menu was very conscientious when it came to the truth of their statements. They probably interrogated the chef to know what dishes he intended to cook and serve, and probably had the menu checked by the restaurant's personnel, to ensure that he did not make any error or false statement. We can also imagine that the person who composed a flattering resume made sure that everything he claimed on his resume was true in some sense. Thus, it does not seem that Moberger's account is able to capture such instances of bullshit.

More generally, it seems to me that attempts at defining bullshit by reference to some defects in the process of inquiry that led someone to endorse or utter a certain claim runs the risk of returning too many false positives. Indeed, imagine a person named John who watches a documentary on TV. In one case, the documentary turns out to contain certain pseudo-scientific claims (e.g. that homeopathy can cure cancer). In another, parallel case, the documentary contains scientific claims (e.g. that homeopathy does not work). In both cases, because John generally trusts documentaries, he accepts the documentary's claims, and then repeats them in a conversation (without pushing the inquiry further). This means that John's attitudes towards his evidence and epistemic conscientiousness is *identical* in both cases. However, though it seems right to say that John spouts bullshit when he claims that "homeopathy can cure cancer", it also seems weird to say that John is bullshitting or propagates bullshit when he claims that "homeopathy does not work". Of course, we could bite this bullet but it seems that most people accept and share claims with the same degree of inquiry and concern for evidence as John, meaning that we should conclude that most of



what we say is bullshit. This would make the concept of bullshit useless as (i) it would encompass too many cases, and (ii) it would make a lot of bullshit simply unrecognizable. Indeed, how would you know whether someone believes in evolutionary theory as a result of a rigorous and conscientious inquiry, or simply because he trusted some documentary without giving it much thought?

### 3. An Alternate, Output-Based Proposal

Thus, it seems that most existing accounts cannot accommodate cases of “truth-tracking bullshit.” Now, all the accounts I discussed so far are “process-based” rather than “output-based”—meaning that they propose to define bullshit not based on the properties of the statements that are produced, but rather on the process that lead the agents to produce these statements. Maybe an “output-based” approach, focusing on the properties of the statements themselves, would be more successful.

This distinction between two approaches to bullshit comes from Cohen (2006). He bases this distinction on the Oxford English Dictionary, which gives the following two definitions of bullshit: (1) “nonsense rubbish,” and (2) “trivial insincere talk or writing” (Cohen 2006: 120). Whether something is bullshit in the first sense depends on the properties of what is produced, while whether something is bullshit in the second sense depends on the intentions, mental states and attitudes of the person who produces it. In this view, bullshit in the first sense is a noun that emphasizes the worthlessness of bullshit, while bullshit in the second sense is a verb that focuses on the process of bullshitting (Cohen 2006: 120–121). The first sense is *output-based*, while the second is *process-based*.

According to Cohen, Frankfurt’s process-based account is insufficient and should be completed with an output-based account of bullshit. He takes as an example his own experience with French Marxism, and more precisely with the Althusserian school:

The ideas that Althusserians generated ... possessed a surface allure, but it seemed often impossible to determine whether or not the theses in which those ideas figured were true, and, at other times, those theses seemed capable of just two interpretations: on one of them, they were true but uninteresting, and, on the other, they were quite interesting, but obviously false. (Failure to distinguish those opposed interpretations produces an illusory impression of interesting truth). (Cohen 2006:118)

For Cohen, such obscure philosophical texts are paradigmatic cases of bullshit, but Frankfurt’s account of bullshit does not really capture them: it is not clear

that people who produce such sentences are not concerned with truth. In fact, Cohen vouches for them, claiming that at least some of the philosophers who produce these kind of sentences are “honest thinkers.”

However, Cohen does not take such examples to constitute a counter-example to Frankfurt’s analysis of bullshit. Rather, Cohen concludes that we should distinguish between two different types of bullshit: bullshit in ordinary life, which is best captured by Frankfurt’s approach, and bullshit in academic work, which requires another kind of approach. To capture the second kind of bullshit, Cohen proposes that a statement or text is bullshit when it is *nonsense*, that is when it is

by nature *unclarifiable*, discourse, that is, that is not only obscure, but which cannot be rendered unobscure, where any apparent success in rendering it unobscure creates something that isn’t recognizable as a version of what was said. (130)

One objection often leveraged against Cohen’s definition of bullshit is that it is very narrow and that a lot of what we consider bullshit still has meaning. In fact, it is not even clear that Cohen’s proposal can account for his experience with Althusserian texts: according to him, such texts are clarifiable but, depending on how they are clarified, their claims are either obviously false or trivially true. As such, they correspond to what Dennett (2009) called *deepities*: misleading statements that can be interpreted in two ways: a true but trivial one, and a surprising but obviously false one.

However, Cohen himself admits that “defects other than unclarifiable unclarity can suffice to stigmatize a text as bullshit” (Cohen 2006: 131). He hints at other properties that might flag a statement as bullshit, such as “arguments that are grossly deficient, either in logic or sensitivity to empirical evidence”, or “irretrievably speculative comments.” As such, it might be possible to broaden Cohen’s proposal and improve on it with a more comprehensive output-based account of bullshit.

But where to begin? This is where I must emphasize the second purpose of the examples I presented in introduction: as I see it, they draw our attention to an important feature of bullshit—namely that bullshit is not necessarily meaningless, false, or lacking in evidence or respect for truth, but that it can always be *deflated*. Indeed, what makes the menu or the flattering resume instances of bullshit is both that they present the truth in an overly flattering light and that this impression will disappear as soon as we begin investigating what they really mean.

The same is true for other examples of bullshit. Lacan’s statement might sound deep and impressive but wondering “what they really mean” will lead us to conclude that they are meaningless. The Althusserian texts described by

Cohen and Dennett's *deepities* might also give us, at first sight, the impression of encapsulating some profound insight, but asking ourselves "what they really mean" will lead us to conclude that they are either false or trivial. Pseudo-sciences promise to give us some understanding of the world we live in, but this impression is misleading. Thus, the common thread between all these cases is the following: bullshit is something that seems to have much value at first sight, but scratch the surface and you will see that it is *empty*. To put it simply: bullshit is just *hot air*.

I propose to formalize this very simple insight in the following way:

*What makes a given claim C bullshit is that (i) though C is presented as or appears at first sight as making an interesting contribution to a certain inquiry,<sup>5</sup> (ii) C would turn out, on closer inspection by a minimally competent inquirer, to make a much less interesting contribution.*

Here, I introduce the notion of "minimally competent inquirer", because we don't want all claims that ultimately fail to make an interesting contribution to count as bullshit. Imagine that you put forward a bold but plausible scientific hypothesis. A colleague then goes to great lengths to test your hypothesis. After years of effort, he finally succeeds in disproving your theory. Would it be fair for him to conclude that you were "bullshitting?" Probably not. Our intuition is rather that bullshit can be proved worthless via a very basic, elementary inquiry. You just have to scratch the surface to see that it's empty.

We can thus define a minimally competent inquirer as someone who is not necessarily a specialist but has a basic comprehension of the terms involved in a bullshit claim and (when applicable) of the theories they play a role in. For example, psychologists Evans and colleagues (2020) designed a "receptivity to scientific bullshit scale," which mixes scientific and pseudoscientific statements such as:

- In all thermal equilibria, if no surface tension is applied nor any refraction imposed upon the object, the capacity of that atomic structure disperses throughout the object.
- In a natural thermodynamic process, the sum of the entropies of the interacting thermodynamic systems increases.

---

5. I must admit that I rely here on an intuitive understanding of "interesting". This is because I think that "interestingness" is a response-dependent property. However, I can try and characterize it in the following way: a contribution is interesting if (i) it answers a question that matters to us, and (ii) answers it in a way that is either surprising or illuminating (i.e. the claim must either bring new, unexpected, or game-changing information, or help us better integrate the information we already had in a coherent explanatory framework). The most interesting claims are thus those which provide new information that sheds new light on past information we had.

In this case, the first sentence is bullshit, while the second is scientifically valid. If you guessed on your own—congratulations!—I was personally unable to tease them apart. However, the point is that most people who have a basic comprehension of the scientific terms involved in these sentences and of the theory they play a role in would certainly be able to. This is what makes the scientific bullshit statement bullshit and not some very speculative but really interesting hypothesis.<sup>6</sup>

In addition to this knowledge, the minimally competent inquirer should also master the basics of formal and informal logic, and be able to spot the most obvious forms of bad reasoning and arguments. They should also have a certain quantity of basic knowledge – meaning that they should know what most people in their epistemic community know.<sup>7</sup> In some cases in which the target claim contains no technical term, then the minimally competent inquirer is just a person with this basic knowledge and a decent mastery of logic and reasoning.

The phrase “on closer inspection” should also be clarified. This includes at least two steps. The first (*conceptual clarification*) is the willingness to clarify what exactly is said, to ask oneself the question: “what does it really mean?” The second (*fact-checking*) is the willingness to seek easily accessible information to verify whether the claims that are made are (i) non-trivial, and (ii) true or at least based on reasonable evidence. As we saw, some claims, such as Lacan’s claim, do not pass the first step. Others, such as the main tenets of homeopathy, might pass the first, but not the second.

With that in mind, I can thus summarize this new proposal in the following way: bullshit is what seems or purports to make an interesting contribution to a certain inquiry but can be identified as failing to do so under closer inspection by a minimally competent inquirer. But there are many ways in which a claim C can fulfill these conditions. Here are some of them:

- C is meaningless,
- C has meaning but is obviously false,
- C has meaning and is true but is trivial or uninteresting,

---

6. In the same way, most non-specialists would be unable to detect the misuses of scientific concepts and theories listed by Sokal & Bricmont (1997), but what makes these examples bullshit is that anyone with a basic mastery of these concepts and theories would see that there is something fishy.

7. This is why the very same claim (e.g. the Earth is at the center of the solar system) can count as bullshit in a certain context (e.g. 21<sup>st</sup> century Europe), but not in another one (e.g. 15<sup>th</sup> century Europe). Moberger (2020: 598) takes this context-dependency as an argument for the process-based accounts of bullshit. However, I don’t think that this shows that the producer’s intentions and dispositions are relevant when assessing whether a claim is bullshit: rather, it just shows that the social context in which a claim is made should be taken into account when assessing this claim.

- C has meaning, is true, and is not trivial, but there is no way that it can be grounded in evidence,
- C has meaning, is true, is not trivial, is plausibly grounded in sufficient evidence, but is irrelevant to the context of the discussion.

Statements that purport or appear to make an interesting contribution but ultimately turn out to be meaningless are thus only one sort of case of bullshit statements, but are probably the most paradigmatic ones, as closer inspection will ultimately prove them 'empty' and 'void.' Note that there are several ways in which a statement can be meaningless. Some might contain meaningless terms or pseudo-words, such as "that's a nice Tnetennba" or "the geodesic biocompensator was tested using a Teslameter 05/40 and the Voltcraft Vc 960" (Milgram 2021). Others might contain real, meaningful words but combine them in a way that fails to make sense, such as "the sum of the angles of a triangle is equal to the color red" or "quadruplicity drinks procrastination." However, in line with Erwin (1970: 161), I propose that meaningless statements all share one common property: "we cannot understand what it would be like for them to be true". More precisely, I propose to cash out this intuition in the following terms: a statement is meaningless when it is impossible for competent speakers to conceive a possible world or state of affairs that would make this statement true.<sup>8</sup>

#### 4. Arguments for the Proposal

Now that I have put forward my own account of bullshit, I would like to list a few arguments in its favor.

(1) First, I think this definition allows for the integration of most of the examples of bullshit put forward in the literature. Because it includes meaningless sentences, it accounts for cases of "meaningless" and "unclarifiable" bullshit, such as Lacan's claim (Cohen 2006). Because it includes claims that are meaningful, but notoriously false, it also includes most cases of pseudoscience, such as homeopathy, astrology, flat earth theories, and other kinds of conspiracy theories (Johnson 2010; Dieguez 2018; Moberger 2020). Because it includes claims that are true but that turn out to be trivial or uninteresting after clarification, it includes "deepities" (Dennett 2009). Because it includes claims that are true

---

8. More precisely, using Chalmers' terminology (2002), I would say that a statement is meaningless when there is no *prima facie* positively conceivable world that would make this statement true. The emphasis on *prima facie* is important to bypass one difficulty of Erwin's proposal to define meaningless statements as statements that are "*a priori* false." This definition is too stringent because it would exclude as meaningless mathematical statements that can be demonstrated false but are not obviously so.

and non-trivial, but that are made without proper evidence, it also allows us to include the cases of “irretrievably speculative comment” put forward by Stokke and Fallis (2007). And because it includes claims that are not relevant to the context of discussion, it includes cases of evasive bullshit (Carson 2016).

One question, however is whether my account can accommodate two further types of bullshit identified by Moberger (2020): *obscurantist pseudophilosophy* (the kind of discourse produced by people in humanities and social sciences when they engage with philosophical issues without awareness of the relevant distinctions and arguments), and *scientific pseudophilosophy* (the kind of discourse produced by people with background in natural sciences when they engage with philosophical issues without awareness of the relevant distinctions and arguments). What makes such discourses bullshit according to Moberger is that their author’s ignorance of the relevant philosophical literature shows that they do not seriously engage in the relevant epistemic inquiry (i.e., they are not “conscientious” enough). However, I disagree: Let’s imagine that a scientist ponders over a serious philosophical topic in isolation, being completely ignorant of the relevant philosophical literature, but produces insightful theories and arguments regarding this topic. Should we still label his proposals as “bullshit”? I don’t think so. Rather, what makes obscurantist and scientific pseudophilosophy bullshit is that anyone with a basic knowledge of the relevant philosophical literature can see that such claims are conceptually confused, trivially false, or insufficiently warranted by the available evidence (often because the existence of a large philosophical literature discussing these topics act as a defeater for most of the simplistic views and arguments presented by pseudophilosophers).

(2) Second, I think that my account allows for a satisfactory elucidation of the relationship between lies and bullshit. On the one hand, certain accounts (such as Frankfurt’s account or Stokke and Fallis’ account) draw a sharp distinction between the two: if a claim is a lie, it cannot also be bullshit. On the other hand, most prototypical cases of bullshit in ordinary language are also cases of lies. For example, one famous best-seller about bullshit opens with the following claim:

We live in an era of unprecedented bullshit production (...) Never in history have so many people uttered statements that they know to be untrue. (Penny 2005:6)

Thus, it seems that bullshit in the ordinary sense somewhat intersects with lies. My account explains how some lies can be bullshit, without counting all lies as bullshit: a lie is bullshit when it purports to convey very important and interesting information. But most lies are not of this kind, especially when people who lie do so to keep a low profile. For example, a cheating husband who is asked where he was at noon should opt for a mundane answer such as “I was having



lunch with Bob from the Accounting Department”, rather than “I was having lunch with Barack Obama.”

(3) Additionally, my account of bullshit is more in line with the way bullshit is conceptualized and operationalized in other academic disciplines. For example, anthropologist David Graeber (2018) defines “bullshit jobs” in the following way:

a bullshit job is a form of paid employment that is so completely pointless, unnecessary, or pernicious that even the employee cannot justify its existence even though, as part of the conditions of employment, the employee feels obliged to pretend that this is not the case. (9–10)

In other words, a bullshit job is a job that gives the impression of or is touted as having an important impact while, in fact, making no contribution whatsoever. This is in line with my definition of bullshit and allows for a unification of the category of “bullshit” across objects and disciplines.

Moreover, in the past ten years, there has been a growing body of research on the psychology of “pseudo-profound bullshit”. In their seminal paper, which gave birth to this research program, Pennycook and colleagues (2015) characterize bullshit in the following way: “bullshit, in contrast to mere nonsense, is something that implies but does not contain adequate meaning or truth”. As can be seen, this is an *output-based* rather than a *process-based* definition of bullshit. But more important is the method Pennycook and colleagues use to produce their list of bullshit statements: they use statements that have been generated at random by non-human programs such as “The New Age Bullshit Generator” or “Wisdom of Chopra”. However, it is hard to explain why such statements (or, more generally, statements generated by large-language models) can be bullshit on the basis of most *process-based* accounts, for which the indifference of the locutor towards truth is what makes bullshit bullshit.

(4) The case of artificially-generated statements lead us directly to a fourth argument already advanced by Cohen (2006), which is that, to identify bullshit, we rely on the properties of what is produced, rather on inferences about the motives and dispositions of the person who produced it:

Think of attempts to vindicate Heidegger, or Hegel. The way to show that they weren't bullshitters is not by showing that they cared about the truth, but by showing that what they said, resourcefully construed, makes sense. (130)

Conversely, when we try to demonstrate that something is bullshit, we do so by showing that it is meaningless, obviously false, or much more trivial than it

appears at first sight. We do not appeal to its author's intention or mental states. This is more in line with *output-based* accounts than *process-based* accounts.<sup>9</sup>

(5) My account can explain the *bullshit* metaphor: why is it always bullshit or horseshit, but not flyshit or mouseshit? It's because the excrement of bulls and horses are *impressively huge* (when compared to the ones of flies and mice) – but big shit is still shit. They might seem impressive at first sight but are worthless on closer chemical analyses – which is basically what my analysis of bullshit is getting at.<sup>10</sup>

(6) Finally, one last argument for my account is simply that it provides a simple explanation for the attractiveness of *process-based* accounts. Indeed, once we realize that a claim is obviously meaningless, false or trivial, it is hard to accept that the person who made that claim could have done so while being genuinely concerned about the truth. Indeed, that would mean that this person is somewhat stupid, or at least much less clever than we are. For people who are charitable, it is thus difficult not to jump from the mere perception that a claim is bullshit to the conclusion that the utterer did not really care about truth. In other words, I think that defenders of *process-based* accounts are too charitable.

I'm not.

## 5. A Process-Based Account of Bullshitting

Still, one might argue that my account of bullshit is incomplete to the extent that it lacks an account of *bullshitting*. A simple solution would be to define “bullshitting” as “intentionally producing bullshit”. However, I don't think that would be satisfactory, for I think it is possible to be a bullshitter without producing bullshit statements. This is the case, I think of many pseudo-experts who pass as experts: they often do so by memorizing a few facts or anecdotes that most people with a few years of training in a given field would know, and then repeat these facts and anecdotes in front of a non-expert audience to convey and give the impression that they are indeed experts (Fuhrer et al. 2021). To many experts, this would appear as “bullshitting” the audience— even though nothing that the pseudoexpert says is “bullshit” in a strict sense (the information might be trivial to the expert, but it is not to the general audience the pseudoexpert is trying to seduce).

---

9. For example, Dalton (2016) objected to Pennycook and colleagues by arguing that the randomly generated statements they used were not really bullshit, because they “could have provided glimpses of insight and wisdom to the subjects” (independently from the process that generated them).

10. But why not “elephant shit”, then? In fact, subtle distinctions between “chickenshit”, “bullshit”, and “elephantshit” have been theorized by psychiatrist Fritz Perls (see Perls 1969).

But maybe a clearer example of “bullshitting without bullshit” might be found in cases of “moral grandstanding” (Tosi & Warmke 2016: 199). Imagine a public figure that takes very harsh (but correct or plausible) moral positions against people who violate a certain moral norm. At some point, it is discovered that this public figure never really cared about this moral norm (or moral norms in general) and that they engaged in moral condemnation only to improve their public image. It would be natural to think that this public figure was “bullshitting” the audience the whole time. However, none of their positions or moral statements need to be bullshit in the sense I defined it.<sup>11</sup>

A last, but even more striking example of “bullshitting without bullshit”, can be found in advertising. Throughout the literature on bullshit, advertising is considered a prototypical example of bullshit. However, a lot of advertisements do not even make claims (Johnson 2010). Good examples can be found in advertisements for perfumes (such as advertisements by Chanel or L’Oréal), which are often sequences of beautiful but puzzling imagery without any words spoken. These advertisements make no claim—as such they cannot be bullshit according to my account (and most accounts I reviewed). Still, they make us feel like we are being bullshited. What gives?

The answer might lie in the fact that, though they do not make any claim, such advertisements seek to produce a certain effect: they want us to associate positive feelings and ideas to their products. This is in line with several suggestions according to which bullshit is more interested in producing effects than communicating ideas. For example, Neumann (2006: 203) writes that: “Bullshit is a certain kind of speech, intended to distract or obfuscate in a general way, in order to achieve a desired effect – often one that is nonrational and emotional, where emotions become reasons for a course of action”. Similarly, Preti (2006: 21) writes: “what matters to bullshit is that it should make something matter to us”. And, speaking of “clickbait” headlines, Bergstrom and West (2020:23) conclude that: “most successful headlines don’t convey facts, *they promise you an emotional experience*”.

Thus one could be tempted to define bullshitting in the following way:

*X is bullshitting when X engages in a communicative act but is more concerned about the general (affective) impression their act will have on a given target T than about the particular propositions they will get T to endorse.*

Such an “impression-based” definition might account for Frankfurt’s intuition that bullshit “is never finely crafted, that in the making of it there is never the

---

11. To *moral grandstanding*, we could add *epistemic grandstanding*. Indeed, since the success of Frankfurt’s book on bullshit, I have come across a handful of talks and conferences about “bullshit” that seemed less concerned with making a theoretical contribution to our understanding of the phenomenon, and more with signaling their authors’ commitment to the “values of knowledge”.

meticulously attentive concern with detail” (Frankfurt 2005: 21), while “the liar is inescapably concerned with truth-values” (51). Indeed, while the prototypical liar is concerned with implanting a particular belief in a particular proposition (e.g. that he was out of town when a certain murder was committed), the prototypical bullshitter is more concerned with leaving a general impression (e.g. that he’s competent in his field).

However, despite its advantages, this definition of bullshitting might still be too broad to be adequate. Imagine someone who is composing a CV and a motivation letter to apply to a job but has enough accomplishments and genuine motivation to avoid writing bullshit in the sense I defined it. Still, his main goal when engaging in those actions is to leave a positive impression on potential employers. But it would be strange to conclude that this person is “bullshitting”.

I think we can avoid such strange conclusions by introducing into my definition of “bullshitting” something I left out of my definition of “bullshit”: a certain modicum of dishonesty. Thus, we could say that the bullshitter knows that the impression he wants to convey is not clearly warranted. Taking into account that one can bullshit not only to give positive impressions, but also to minimize bad impressions (for example when Putin announces “a special military operation aimed at demilitarizing and denazifying Ukraine” rather than crudely speaking of “war” and “invasion”), or to give terribly negative impressions (for example of political opponents), I can offer a more appropriate definition of “bullshitting” along these lines:

*X is bullshitting when X engages in a communicative act C but is more concerned with the general (affective) impression their act will have on a given target T, than about the particular propositions they will get T to endorse AND X is aware that X has no good reason to think that the impression X is trying to convey is warranted by reasons presented in or hinted at by his communicative act C.*

This definition allows both “bullshit without bullshitting” (by reintroducing the idea that bullshitting involves some kind of dishonesty, while bullshit doesn’t) and “bullshitting without bullshit”. Still, it keeps a unity between bullshit and bullshitting by having both revolve around the idea of “something that is more impressive at first sight than it really is or purported to be”.

## 6. The Problem with Bullshit

So far, I have mainly motivated my account of bullshit on the basis that it allows me to accommodate paradigmatic instances of bullshit that are left aside by other existing accounts. However, one might object that the aim of a philosophical

account of bullshit should not be to capture our ordinary linguistic intuitions about what counts as “bullshit”, but to engineer a concept of “bullshit” that is able to make some philosophical work by (i) identifying a certain phenomenon that is worth investigating, (ii) furthering our understanding of this phenomena, and (iii) furthering our understanding of related phenomena. For example, Stokke and Fallis (2017) are interested in bullshit because it allows them to explore the varieties of non-alethic speech beyond lying, while Moberger (2020) tries to put forward a concept of bullshit that would allow him to better understand pseudoscience.

I actually agree with this idea: my inquiry is not an inquiry into ordinary language. I think the goal of a proper philosophical account of bullshit is to provide theoretical insight into phenomena themselves, not on how people use the word “bullshit”. However, I will argue that the account I provided in this paper is better suited to fulfill some of the expectations that we can have about a philosophical account of bullshit.

Indeed, most of the philosophical literature on bullshit touts bullshit as a formidable threat to truth. For example, Frankfurt writes that “bullshit is a greater enemy of the truth than lies are” (2005: 61). Thus, we would like an adequate account of bullshit to explain *why* bullshit is such a threat, and *how* it spreads, so we can properly understand and react to it. However, I don’t think that most philosophical accounts of bullshit succeed in this regard. (Though they might succeed with respect to other criteria.)

### 6.1. *Explaining the Spread of Bullshit*

*Process-based* accounts of bullshit mainly define bullshit as the *symptom* of something else, be it indifference to truth or lack of conscientious inquiry. As such, the first defect of such approaches is that they make bullshit explanatorily impotent: bullshit by itself explains nothing. Rather, it has to be explained by a greater threat of which bullshit is a simple emanation or manifestation (indifference to truth, or post-truth).

Moreover, because they only define bullshit by reference to the process that gave birth to it, these approaches cannot really say what’s wrong with bullshit itself (but only with people who produce it). Of course, one can always say that claims that are made without regard for truth are more likely to be false or misleading, but such features of bullshit are incidental in such accounts, making the problem caused by the spread of bullshit more indirect.

Finally, I think process-based accounts of bullshit have the disadvantage of explaining bullshit mainly in terms of individual failings: if bullshit is (by its very definition) the product of the locutor’s lack of concern towards truth, or

lack of epistemic conscientiousness, then explanations by environmental forces and social contexts are only secondary to explanations in terms of individual epistemic virtues.

However, I think this perspective is erroneous because, in the relevant contexts, anyone is susceptible to giving in to the temptation of spewing bullshit— even an epistemically virtuous person like my reader! Let's take a simple example: letters of reference. I think that most people would agree that reference letters constitute a distinctive type of bullshit (Richardson 2006). At some point, it was typical for us Europeans to mock the hyperbolic style of US reference letters, but we are in the process of quickly catching up – especially when our students want to apply to US universities. Are all the professors writing these reference letters indifferent to the truth? Maybe at the very moment they write these letters, but surely not most of the time. The truth is that most professors are compelled to write hyperbolic bullshit to give their students a chance of competing – even though they fully realize it's bullshit and generally despise bullshit.

To explain why all these professors finally yield to the power of bullshit, it is useless to invoke their indifference to truth or inquiry. Rather, it is best to point to the fact that bullshit is generally more interesting than plain truth and that, in a context in which the most attention-grabbing reference letters tend to win, bullshit is at an advantage. This is an arms race that forces even the most epistemically virtuous individuals to produce bullshit.<sup>12</sup> Similarly, the same professors who deplore the pervasiveness of bullshit will still produce some when the time comes to fill the “novelty of the proposal” and “social implications” sections of their next grant proposal.

In fact, focusing on the fact that bullshit is, by its very nature, more interesting and seductive than other types of discourse might explain why bullshit flourishes in certain areas more than others. Simply put: we would expect bullshit to triumph in areas (i) in which people's success depends on the positive impression others have of them, and (ii) in which impressions are mostly dependent on discourse (a competitive sportsman does need to talk to showcase his qualities). Unsurprisingly, bullshit is abundant on the job market (CVs and reference letters), in advertising (Johnson 2010), in politics (Gibbons, in press), and in philosophy (Cohen 2012).

Thus, compared to *process-based* accounts of bullshit, the *output-based* account I proposed in this paper is best suited to explain why (and where) bullshit is pervasive.

---

12. One could object that the perfectly epistemically virtuous professors should simply sacrifice their students' opportunities to their own epistemic integrity. However, epistemic concerns should sometimes be bypassed by moral concerns.



## 6.2. Why is Bullshit a Threat?

I will also argue that *output-based* accounts can more readily explain what's wrong with bullshit. To begin with, some epistemic defects of bullshit are embedded in the very definition I offered: bullshit is meaningless, false, trivial, ungrounded or irrelevant. This contrasts with the bullshit of process-based accounts, which is the product of epistemically problematic dispositions, but can still be true, interesting and relevant.

But, more importantly, my account of bullshit can explain why bullshit is a threat to truth and why it is not only a *sign*, but also a *source* of indifference towards truth. Indeed, at the core of my account lies the idea that bullshit is seductive, intellectually attractive – but also that it falls flat upon closer inspection. This means that bullshit provides us with a certain satisfaction, but that this satisfaction is dependent on our refraining from investigating further. As such, bullshit might discourage us from inquiring into the truth, as well as foster hostility towards those who would promote such inquiry.

Frankfurt (2005) famously claimed that “people do tend to be more tolerant of bullshit than of lies” (50)—an observation that seems to be supported by recent empirical results (Petrocelli et al. 2023). Reasons for this tolerance might integrate the fact that lies are false by definition, while bullshit can sometimes be true (Petrocelli et al. 2023), or the fact that bullshit can be outright funny (Kimbrough 2006). But, I think that most of our indulgence is best explained by the fact that we find certain pleasant qualities in bullshit.

Let's for example go back to Pennycook and colleagues' “pseudo-profound bullshit.” Further studies have found that the more people find these randomly-generated statements profound, the more they report being positively “moved” or “touched” by them (Cova, Deonna & Sander 2018; Cova & Boudesseul 2023). Similarly, recent studies suggest that “positive” pseudo-profound bullshit is more likely to be rated as profound than “negative” bullshit (Altay et al., 2023). This suggests that people who find such statements profound also take a certain form of pleasure and satisfaction in engaging with them. Let's now assume that this satisfaction would go away if they began assessing them in a more critical way. Wouldn't that deter one from engaging in such inquiry?

More generally, bullshit can be comforting, which might make people who “call bullshit” look like killjoys or party-poopers. For example, Aberdein asks us to imagine the case of a person “who tells a critically injured person that “Help is on its way”, despite having no idea whether this was true, because [he] was hoping for the best and did not wish to needlessly demoralize someone clinging to life” (2006, 167). Is that bullshit? Aberdein answers negatively but my own definition of bullshit leads to the following conclusion: that is definitely bullshit. Telling people you love that you will “love them no matter what” is also bullshit

(as you would probably stop loving them if they were suddenly revealed to be serial killers or the type of people who regularly submit papers but never accept to review any). So is telling your friends, children or students that “they can do whatever they put their mind to.” But we engage regularly in this sort of comforting bullshit, and would frown on people who would oppose such behavior.

Of course, this is just a particular case: not all bullshit is comforting (let’s think about conspiracy theories). However, there is one positive quality it always has, or rather always seems to have: by my own definition, it always seems interesting. And, I surmise, we can enjoy interestingness for itself, independently from truth (even though having in mind the falsity of a claim typically prevents us from really enjoying its interestingness).

Indeed, though Aristotle claimed that human beings seek knowledge, it is clear that we do not seek just *any* form of knowledge. There are a lot of things we do not care to know, such as the exact number of hairs on the head of our neighbor or the exact number of characters in the philosophy paper we are currently reading. Generally, our epistemic interests are directed towards a subset of truths: *interesting* truths. We can imagine that there are good evolutionary reasons for this, but it means that our drive for knowledge and truth is mingled with a drive for interestingness. And that, on some occasions, these drives come apart and we might be able to enjoy interestingness for itself. As Whitehead once claimed: “it is more important that a proposition be interesting than that it be true” (see Stace 1944, 233).

There is a certain intellectual satisfaction in engaging with interesting claims and ideas. As Stace (1944: 235) writes: “Our interest is aroused by the bare fact of newness. A new idea is like a new sauce. It tickles the intellectual palate as a new sauce tickles the physical palate. We value this new sensation.” However, bullshit is a case in which this satisfaction is divorced from truth and plausibility, which means that fully appreciating the “interesting” part requires leaving aside concerns about “truth” as a boring afterthought. This is, at least, how Cohen (2012) explains why bullshit is so pervasive in French intellectual life: according to him, French intellectuals write mainly for a general audience, and “audience will read philosophy only if it is *interesting*, and being interested in interestingness is quite different from being interested in truth” (110). In the same way people who find comfort in bullshit might become averse to inquiring about truth and people who promote this kind of inquiry, people who relish in the interestingness of ideas might lose most of their concern for truth and might come to shun the boring fact-checkers who threaten to spoil their fun.

Thus, because bullshit provides certain emotional and intellectual satisfactions, it constitutes a case in which our interest in truth conflicts with other interests (including epistemic interests such as our interest in what is interesting). Because bullshit is also fragile, these competing interests discourage

serious inquiry into the truth of bullshit, and foster hostility towards those who would be willing to “call bullshit.” As such, bullshit itself, by its seductiveness, might contribute to cultivating a defiance towards reason and rational inquiry.

What can be done about it? I honestly don't know. However, I (somewhat) hope that the account of bullshit I put forward in this paper might help us to understand how bullshit spreads and how it can be countered. Though I must say that I highly doubt that this will be done by multiplying incantations to epistemic virtues and individual epistemic responsibility. I think that bullshit is more of a collective issue, and that one way to fight it might be to collectively rehabilitate being boring.

## Acknowledgments

This work was supported by an Eccellenza Professorial Fellowship of the Swiss National Science Foundation (“Eudaimonic emotions and the (meta-)philosophy of well-being”). I would like to thank Jay Richardson for carefully proofreading the manuscript, Phersv for introducing me to Sebastian Marx's show, Kenzo Nera for the examples of meaningless statements, and Mark Steen for his suggestions on the final manuscript. I initially planned to dedicate this manuscript to Carole and HN, whose preference for interestingness over truth sparked most of the reflections I developed here. However, given the recent passing of Harry G. Frankfurt, who was an inspiration in so many philosophical domains, I finally decided to dedicate to him this modest contribution to a research program he initiated.

## References

- Aberdein, Andrew (2006). Raising the tone: Definition, bullshit, and the definition of bullshit. In Gary L. Hardcastle and George A. Reisch (Eds.), *Bullshit and Philosophy* (151–169). Open Court.
- Aberkane, Idriss (2021). Allocution publique pour le retour à la liberté devant les Nations Unies à Genève. Retrieved from: <https://www.youtube.com/watch?v=E8jppqND6GyQ>
- Altay, S., Majima, Y., & Mercier, H. (2023). Happy thoughts: The role of communion in accepting and sharing (mis) beliefs. *British Journal of Social Psychology*, 62(4), 1672–1692.
- Carson, Thomas L. (2016). Frankfurt and Cohen on bullshit, bullshiting, deception, lying, and concern with the truth of what one says. *Pragmatics & Cognition*, 23(1), 53–67.
- Chalmers, David J. (2002). Does Conceivability Entail Possibility? In T. Gendler and J. Hawthorne (Eds.), *Conceivability and Possibility* (145–200). Oxford University Press.
- Cohen, Gerald A. (2012). Complete bullshit. In Michael Otsuka (Ed.), *Finding Oneself in the Other* (94–114). Princeton University Press.

- Cova, Florian and Jordane Boudesseul (2023). A validation and comparison of three measures of participants' disposition to feel moved (introducing the Geneva Sentimentality Scale). *Cognition and Emotion*, 37(5), 908–926
- Cova, Florian, Julien Deonna, and David Sander (2018). "That's deep!": The role of being moved and feelings of profundity in the appreciation of serious narratives. In Donald R. Wehrs and Thomas Blake (Eds.), *The Palgrave Handbook of Affect Studies and Textual Criticism* (347–369). Palgrave Macmillan.
- Dalton, Craig (2016). Bullshit for you; transcendence for me. A commentary on "On the reception and detection of pseudo-profound bullshit". *Judgment and Decision-Making*, 11(1), 121–122.
- Dennett, Daniel (2009). The Evolution of Confusion. Conference given at the 2009 AAI Convention. Retrieved from: [https://www.youtube.com/watch?v=D\\_gw8JougLQ](https://www.youtube.com/watch?v=D_gw8JougLQ)
- Dieguez, Sebastian (2018). *Total bullshit! Au cœur de la post-vérité*. Presses Universitaires de France.
- Erwin, Edward (1970). *The Concept of Meaninglessness*. JHU Press.
- Evans, Anthony, Willem Slegers, and Žan Mlakar (2020). Individual differences in receptivity to scientific bullshit. *Judgment and Decision Making*, 15(3), 401–412.
- Frankfurt, Harry G. (2005). *On Bullshit*. Princeton University Press.
- Fuhrer, Joffrey, Florian Cova, Nicolas Gauvrit, and Sebastian Dieguez (2021). Pseudo-expertise: A Conceptual and Theoretical Analysis. *Frontiers in Psychology*, 12, 5214.
- Gibbons, Adam F. (in press). Bullshit in Politics Pays. *Episteme*.
- Graeber, David (2018). *Bullshit Jobs*. Simon & Schuster.
- Johnson, Andrew (2010). A New Take on Deceptive Advertising: Beyond Frankfurt's Analysis of 'BS'. *Business and Professional Ethics Journal*, 29(1/4), 5–32.
- Kimbrough, Scott (2006). On letting it slide. In Gary L. Hardcastle and George A. Reisch (Eds.), *Bullshit and Philosophy* (3–18). Open Court.
- Meibauer, Jörg (2016). Aspects of a theory of bullshit. *Pragmatics & Cognition*, 23(1), 68–91.
- Milgram, G. (2021). Le business anti-ondes : ENQUÊTE et CANULAR. Retrieved from: <https://youtu.be/ZQ5UYMvMN7A>
- Moberger, Victor (2020). Bullshit, Pseudoscience and Pseudophilosophy. *Theoria*, 86(5), 595–611.
- Neumann, Vanessa (2006). Political bullshit and the Stoic story of the self. In Gary L. Hardcastle and George A. Reisch (Eds.), *Bullshit and Philosophy* (203–213). Open Court.
- Penny, Laura (2005). *Your Call Is Important to Us: The Truth About Bullshit*. Emblem Editions.
- Pennycook, Gordon, James A. Cheyne, Nathaniel Barr, Derek J. Koehler, and Jonathan A. Fugelsang (2015). On the reception and detection of pseudo-profound bullshit. *Judgment and Decision-Making*, 10(6), 549–563.
- Perls, Fritz (1969). *In and Out the Garbage Pail*. Gestalt Journal Press.
- Petrocelli, John V., Haley E. Silverman, and Samantha X. Shang (2023). Social perception and influence of lies vs. bullshit: A test of the insidious bullshit hypothesis. *Current Psychology*, 42, 9609–9617.
- Preti, Consuelo (2006). A defense of common sense. In Gary L. Hardcastle & George A. Reisch (Eds.), *Bullshit and Philosophy* (19–32). Open Court.
- Reisch, George A. (2006). The pragmatics of bullshit, intelligently designed. In Gary L. Hardcastle and George A. Reisch (Eds.), *Bullshit and Philosophy* (33–48). Open Court.

- Richardson, Alan (2006). Performing bullshit and the post-sincere condition. In Gary L. Hardcastle and George A. Reisch (Eds.), *Bullshit and Philosophy* (83–97). Open Court.
- Roberts, Craige (2012). Information Structure in Discourse: Toward an Integrated Formal Theory of Pragmatics. *Semantics & Pragmatics*, 5, 1–69.
- Sokal, Alan and Jean Bricmont (1997). *Impostures intellectuelles*. Odile Jacob.
- Stace, Walter T. (1944). Interestingness. *Philosophy*, 19(74), 233–241.
- Stokke, Andreas and Don Fallis (2017). Bullshitting, Lying, and Indifference toward Truth. *Ergo an Open Access Journal of Philosophy*, 4, 277–309.
- Tosi, Justin and Brandon Warmke (2016). Moral Grandstanding. *Philosophy & Public Affairs*, 44(3), 197–217.
- Van Rillaer, Jacques (2019). *Freud & Lacan, des charlatans? Faits et légendes de la psychanalyse*. Mardaga.
- West, Jevin D. and Carl T. Bergstrom (2020). *Calling Bullshit: The Art of Skepticism in a Data-Driven World*. Penguin UK.