

Adrian Briciu*

Bullshit, trust, and evidence

<https://doi.org/10.1515/ip-2021-5003>

Abstract: It has become almost a cliché to say that we live in a post-truth world; that people of all trades speak with an indifference to truth. Speaking with an indifference to how things really are is famously regarded by Harry Frankfurt as the essence of bullshit. This paper aims to contribute to the philosophical and theoretical pragmatics discussion of bullshit. The aim of the paper is to offer a new theoretical analysis of what bullshit is, one that is more encompassing than Frankfurt's original characterization. I part ways with Frankfurt in two points. Firstly, I propose that we should not analyze bullshit in intentional terms (i.e. as indifference). Secondly, I propose that we should not analyze it in relation to truth. Roughly put, I propose that bullshit is best characterized as speaking with carelessness toward the evidence for one's conversational contribution. I bring forward, in the third section, a battery of examples that motivate this characterization. Furthermore, I argue that we can analyze speaking with carelessness toward the evidence in Gricean terms as a violation of the second Quality maxim. I argue that the Quality supermaxim, together with its subordinate maxims, demand that the speaker is *truthful* (contributes only what she believes to be true) and *reliable* (has adequate evidence for her contribution). The bullshitter's main fault lies in being an unreliable interlocutor. I further argue that we should interpret what counts as adequate evidence, as stipulated by the second Quality Maxim, in contextualist terms: the subject matter and implicit epistemic standards determine how much evidence one needs in order to have adequate evidence. I contrast this proposed reading with a subjectivist interpretation of what counts as having adequate evidence and show that they give different predictions. Finally, working with a classic distinction, I argue that we should not understand bullshit as a form of deception but rather as a form of misleading speech.

Keywords: bullshit; deceiving versus misleading; neo-gricean pragmatics; pseudoscience; trust

*Corresponding author: Adrian Briciu, Department of Philosophy and Communication Sciences, West University of Timisoara, Timisoara, Romania, E-mail: adrian.briciu@gmail.com

1 Introduction and methodological remarks

In his popular essay, *On Bullshit*¹, Frankfurt gives a characterization of bullshit as a distinctive form of misleading speech, namely as one whereby the speaker speaks with an indifference toward the truth and intentionally hides her indifference from the audience. What Frankfurt characterized is, as noted by Meibauer (2016), a pragmatic phenomenon: a property of certain utterances that can be characterized in pragmatic terms. Over the years, numerous authors have offered amendments, friendly criticisms and even alternatives to Frankfurt's characterization and by now there are many proposals as to what bullshit, as a form of speech, is.

Methodologically, two approaches can be found in the literature. One is to stick with Frankfurt's definition of bullshit and try to render it in pragmatic terms. This is the line advocated by Dynell (2018). The other approach, taken by those who propose amendments and alternative definitions, is to give a more comprehensive definition of bullshit, one that accounts both for Frankfurt's examples and for more, and render this definition in pragmatic terms². For reasons discussed in Sections 2 and 3 my approach is in line with the latter.

How are we to choose between various competing accounts? This is especially stringent since most criticisms to Frankfurt's account claim that his characterization of the phenomenon is too restrictive. What is, so to speak, the data that should be predicted by such accounts and to which should they be sensitive? Following Frankfurt, it was tacitly accepted that data are people's intuitive judgements about whether particular utterances are bullshit or not. Also following him, most authors have tried to offer a definition of bullshit – that is a characterization in terms of necessary and sufficient conditions. This is a bit problematic because there is no general agreement on what counts as bullshit, or to put it more precisely, people have divergent pre-theoretic intuitions on specific cases. For sure, some of this is due to the fact that “bullshit” is also a term of abuse and is used in regular talk in a rather unconstrained way. But even when these uses are set aside disagreement persists. To give an example, on which I'll say more later, Carson (2010) claims that (some) evasive answers can be characterized as bullshit, and should be covered by an account of bullshit, while others (e.g. Meibauer 2016) claim different intuitions. As long as we have these divergent intuitions on specific cases *some* of these disputes might remain unresolved. Or, a choice between two competing accounts will not be settled by looking at how well they explain and

¹ Initially published in 1986 in *Raritan* and republished as a book in 2005. All references here are to the 2005 edition.

² See for example Cohen (1987/2012), Meibauer (2016), Carson (2010), Davis (2021), Fallis (2014), Stokke (2018a,b) Fallis and Stokke (2017), Wreen (2013).

predict the data, since there is not agreement on what the data is, but also by considerations about other explanatory benefits that such an account will bring. Finally, it might be the case people have divergent intuitions because bullshit is not a completely homogenous phenomenon. This is also suggested by the fact that “bullshit” is polysemous, as it is attested by the OED which lists two distinct entries for the term (1) “Non sense, rubbish” (2) “trivial or insincere way of talking”. Then, the best we can hope for is to give an account of bullshit, as a pragmatic phenomenon, that covers a great deal of cases even if it fails to cover all.

This is precisely the strategy adopted here. I will outline an account of bullshit that illuminates a great deal of cases, including some left out by other accounts, but that doesn’t claim to be exhaustive. In particular, it will not cover what Carson takes to be evasive bullshit, but an alternative explanation of that is given in the last section.

The paper is structured as follows: given its seminal role I’ll present in Section 2 Frankfurt’s main ideas, followed by a short presentation of what I take to be two valid points raised by Cohen’s criticism. In the third section I’ll present a variety of data, some of them new, that motivate and support my account. I will argue that these data suggest that we should not analyze bullshit in relation to truth; and that we should not analyze bullshit in intentional terms. To put it here roughly, I will argue in Section 3 that bullshit is best analyzed as speaking with carelessness toward the evidence for one’s conversational contribution. In Sections 4 and 5 I argue that this characterization can be captured in Gricean terms as a violation of the second Quality maxim, and I offer a non-subjectivist interpretation of the maxim (in contrast with some subjectivist interpretations found in the literature). The rest of Section 5 discusses two alternative accounts and rebuts some putative counterexamples.

2 Frankfurt and Cohen on bullshit

According to Frankfurt, the distinctive feature of bullshit is that the bullshitter speaks with an indifference toward truth and hides her indifference from the audience. “The fact about herself that the bullshitter hides, says Frankfurt, is that the truth-values of her statements are of no central interests to her” (Frankfurt 2005: 55). And, according to him, this indifference is the root cause of bullshitter’s main epistemic fault, namely, that of offering “a description of a certain state of affairs without genuinely submitting to the constraints which the endeavor to provide an accurate representation of reality imposes.” (Frankfurt 2005: 32).

As a paradigmatic example of bullshit, Frankfurt (2005: 16) asks us to consider a politician who on national day bombastically says “our great and blessed

country, whose Founding-Fathers under divine guidance created a new beginning for mankind”. The important feature to note, says Frankfurt, is that the politician is not trying to deceive her audience about American history, but rather she intends to create a certain impression of herself. This is what sets bullshit apart from other forms of misleading speech.

Even if bullshit and lying do not exclude each other, as Frankfurt admits (2005: 23), probably the best way to delineate the former is to contrast it with the latter. As Frankfurt puts it

what bullshit essentially misrepresents is neither the state of affairs to which it refers nor the beliefs of the speaker concerning that state of affairs. Those are what lies misrepresent, by virtue of being false. Since bullshit need not be false, it differs from lies in its misrepresentational intent. The bullshitter may not deceive us, or even intend to do so, either about the facts or about what he takes the facts to be. What he does necessarily attempt to deceive us about is his enterprise. His only indispensably distinctive characteristic is that in a certain way he misrepresents what he is up to (Frankfurt 2005: 54).

This is to say that lying and, I would add, misleading by means of untruthful implicatures are forms of deceit about the facts, or what one takes the facts to be, whereby one deceives her audience about the truth-value of the communicated content. When one lies or invites an untruthful implicature one intends to make her audience believe something to be true which she herself believes to be false. Whereas the liar and the one falsely implicating deliberately promote falsehood, the bullshitter doesn't. What the bullshitter says might be false, but it might as well be true. She doesn't care. “It is just this lack of connection to a concern with truth— this indifference to how things really are— that [Frankfurt] regard[s] as the essence of bullshit.” (Frankfurt 2005, 29). Precisely because the bullshitter's deception concerns her *act* and not necessarily the truth-value of what she communicates, bullshit is not so much a matter of *falsity*, but rather a matter of *phoniness* (Frankfurt 2005: 47).

What does it mean to speak with an indifference toward truth? There are two ways, a broad and a narrow one, in which we can understand this, depending on how we understand the loose term “to speak”. When one speaks with an indifference toward truth, is she indifferent to the truth of strictly what she says, or is she indifferent toward the truth of what she contributes in a conversation by her utterance?

We can interpret speaking with an indifference to truth in a narrow way as speaking with an indifference to the truth of what is said, where *what is said* is to be

understood in Gricean terms³. Then, Frankfurt's characterization of bullshit comes to this: if by utterance *u*, *S* says that *p*, *S* is bullshitting if and only if *S* is indifferent to the truth or falsity of *p*.

Alternatively, we can interpret it in broader terms: when one speaks with an indifference towards truth one is indifferent toward the truth or falsity of what she conversationally contributes. In this interpretation, Frankfurt's proposal amounts to this: a speaker *S* is bullshitting if by an utterance *u*, *S* conversationally contributes that *p* (where *p* is what is said or a conversational implicature carried by *u*) and *S* is indifferent to the truth of *p*.

Although Frankfurt doesn't use this distinction, the existence of *bullshit implicatures* provide a good reason to prefer the broader interpretation. Just as one can say truthful things but mislead the audience by means of untruthful conversational implicatures, so one can bullshit by means of conversational implicatures. One might not be indifferent to the truth or falsity of what one says, indeed might assert a true proposition, but might invite implicatures whose truth are indifferent to her. To see this, consider an adaptation of one of Grice's well-known examples (Grice 1989: 32). Believing Bea to be more knowledgeable about Smith's sentimental life, Ann asks: "Does Smith have a girlfriend these days". Bea has no idea about Smith's personal life but she knows that, for reasons unknown to her, Smith travels frequently to NY. She is indifferent towards his sentimental life and towards Ann's attitudes toward Smith. But she wants to look knowledgeable, so she answers "He has been paying a lot of visits to New York lately" inviting the implicature that Smith has a sentimental partner. If Bea is indifferent toward the truth of her implicature, then this is a case of bullshit implicature.⁴

To sum up, in Frankfurt's characterization bullshit has two distinctive features: one is that the bullshitter is indifferent to the truth-value of what she communicates, and the other is that the bullshitter intends to hide her indifference from the audience. Both features make bullshit intentional. First, indifference is an attitude that one takes and, thus, intentional. Secondly, hiding one's attitude involves an intention to deceive. These are precisely the points targeted by Cohen's criticism.

3 There is considerable debate within philosophy of language when it comes to specifying *what is said*, but for our purposes Grice's general definition will suffice: what is said by an utterance of a sentence is a proposition closely determined by the "conventional meaning of the words (the sentence) uttered" once disambiguation has taken place and the reference of indexical and demonstrative expressions has been fixed (Grice 1989: 25).

4 For bullshit implicatures see Webber (2013: 655–6) and Meibauer (2016: 86). For a discussion of misleading by untruthful implicatures see Saul (2012, chap 4).

Cohen claims⁵ that Frankfurt's account doesn't cover the entire phenomenon of bullshit. He believes that there are what might be called *honest bullshitters* and rejects the two features that Frankfurt identified to be distinctive to bullshit: (a) that the bullshitter is indifferent to the truth of what she contributes and (b) that the bullshitter intends to deceive her audience. He gives the following motivation:

A person who speaks with Frankfurtian indifference to the truth might do so yet *happen* to say something true [...]. And, oppositely, an honest person might read some bullshit that a Frankfurt- bullshitter wrote, believe it to be the truth, and affirm it. When that honest person utters bullshit, she's not showing a disregard for truth. So it is neither necessary nor sufficient for every kind of bullshit that it be produced by one who is informed by indifference to the truth, or, indeed, by any other distinctive intentional state. (Cohen 2012: 104)

According to Cohen, there are at least two forms of bullshit. One is the pedestrian form described by Frankfurt, which is best understood as an activity, and its main feature is indifference to truth. The other type of bullshit is more sophisticated (he claims it occurs mainly in academia), it is best understood as a property of sentences and its main feature is that of being obscurantist, or more exactly, unclarifiably obscure. Cohen doesn't tell us what it means to be clear, but he does tell us what it means to be unclarifiable unclear and he proposes this as a test for detecting bullshit. A statement suffers from unclarifiable unclarity if "adding or subtracting a negation sign from the [statement] makes no difference to its level of plausibility" (Cohen 2012: 105). Unfortunately, only nonsensical sentences pass the test, since only their plausibility is unchanged by adding or subtracting negation. In which case, contrary to Cohen's intentions, this type of bullshit collapses into non-sense. Moreover, as is convincingly argued by Frankfurt in his rejoinder (Frankfurt 2002), many of the examples of unclarifiable unclarity provided by Cohen can be rephrased in a clear manner, although as a result they often turn out to be trite ideas that do not bring anything new to knowledge or understanding.

What motivates Cohen's proposal is the need to account for "the lover of truth who utters what she does not realize is bullshit" (Cohen 2012: 104), that is, to account for the bullshit produced by the honest person who does not show indifference for the truth but has succumbed to the charlatans' charms. Although his positive proposal is, I believe, ultimately unsatisfactory, his criticism raises two important points, namely that there are bullshitters that may care about the truth of their conversational contribution and that bullshitting need not necessarily involve the intention to deceive. I believe that Cohen was right in that the

⁵ His essay was initially published in 1987 and republished in 2012. All references are to the 2012 edition.

phenomenon of bullshit is larger than as conceived by Frankfurt, or to put it differently, in that there are examples of bullshit that cannot be captured by Frankfurt's characterization. In the following section I will provide a battery of such examples. They motivate a departure from Frankfurt's characterization in two points: (1) bullshit should not be characterized in relation to truth; and (2) it should not be characterized in intentional terms. Finally, I will offer a unified account that captures both Frankfurt's cases and those that elude his definition.

3 Departing from Frankfurt

There are good reasons to depart from Frankfurt's analysis of bullshit in terms of speaking with indifference toward truth. We need to depart both from characterizing bullshit in relation to truth, and from characterizing it in intentional terms. The following two examples motivate, I believe, a departure from characterizing it in terms of truth.

3.1 Bullshit: Truth or evidence?

Consider, as a *first example*, someone who is in the habit of making statements for which he has no evidence whatsoever, but nevertheless believes them to be true because "his guts tell him so". Often his statements are contradicted by expert consensus, but this doesn't move him a bit. Sadly, there are real life examples. In a 2018 interview with the Washington Post, when asked for evidence backing some of his claims that were contradicted by specialists, Donald Trump retorted:

I have a gut, and my gut tells me more sometimes than anybody else's brain can ever tell me.
(in Rucker et al. 2018).

In the same interview he also slipped into denialism by rejecting the consensus of 97 percent of climate scientists that climate change is due to human activities motivating his rejection in the following way:

One of the problems that a lot of people like myself, we have very high levels of intelligence but we're not necessarily such believers (i.e. that climate change is due to human activities).
(in Rucker et al. 2018).

Now, given that Trump bluffs blithely we might doubt that he is committed to the truth of what he says, but for sure there are people who match this profile: they are not indifferent toward truth, but are abominably indifferent toward the

evidence for their claims. In other words, although such a person is not indifferent toward the truth of her conversational contribution, there is a sense in which she is indifferent toward the proper evidence for her conversational contribution. I think that such a person can be characterized as bullshitting when she makes such claims. This suggests that one may care about the truth, and nevertheless bullshit if, for example, one fails to observe the evidential requirements of her statements.

As a *second example* that suggests that a proper characterization should not be given in terms of speaking with an indifference toward the truth consider Stokke's wishful thinker.

Lisa is discussing a fishing trip to Lake Mountain View that she has planned to go on with her friends, Vern and Sue. They are all big fans of fishing and have been looking forward to the trip a long time. "I really hope the fishing is good there," Sue says. Lisa has no real evidence about the fishing at Lake Mountain View, and she has no idea what it is like. Still, caught up in the excitement, she exclaims, "The fishing there is outstanding!" (2018a: 267).

I agree with Stokke that many of us will take Lisa to be bullshitting. But she is bullshitting not because she is indifferent to the truth value of her contribution but rather because she is indifferent towards the evidence for her contribution.

These two examples suggest that a proper characterization of bullshit should not be given in terms of indifference toward the truth, but rather as *indifference toward the evidence*. Then we might try a first characterization of bullshit along the following lines: if by utterance u , S conversationally contributes that p , S is bullshitting if and only if S is indifferent to the evidence for the truth or falsity of p .

This is a more encompassing characterization of the phenomenon, than the one in terms of speaking with indifference toward the truth. The latter is merely a particular case of the former. When one is indifferent towards the truth of what she communicates, she is bound to be indifferent towards the evidence for her statements, but not necessarily the other way around. Frankfurt's example of the political orator is a case in point (Frankfurt 2005: 16). If the orator doesn't care whether it is true or not that the founding fathers worked under divine guidance then, *a fortiori*, he doesn't care for the evidence for this claim. But as the above examples suggest, the converse is not the case: one might care about the truth value of her claims but not about the proper evidence for them.

This, I believe, is one step in the right direction. But we also need to give up analyzing bullshit in intentional terms. A few examples motivate this move.

3.2 Bullshit: Indifference or carelessness?

As a *first example*, consider pseudo-scientific statements⁶. Imagine someone who comes to believe ideas long discredited within the scientific community and denies that HIV is the cause of AIDS and repeatedly declares “*HIV does not cause AIDS*”. Secondly, consider a far too common case, unfortunately: a parent who has to decide whether to give her child the MMR vaccine. Suppose that instead of asking her doctor’s opinion she decides to conduct her own ‘research’ over the Internet. As a result of what she reads there she decides not to vaccinate her children and whenever she’s asked about her decision she replies “*Vaccines cause autism*”. Or consider someone who comes to believe that a great deal of modern geology, paleontology, and biology are mistaken and that the universe and all of life came into existence within one week, less than 10,000 years ago. To explain the vast amounts of fossil evidence available for the gradual evolution of species she says “*Only a massive ancient flood could have buried all the fossils fast enough to insure their preservation*”. Finally, consider someone who comes to believe, as a result of a massive and well-coordinated disinformation campaign, that there is no consensus within the scientific and medical communities with respect to the evidence linking cigarettes and cancer. Whenever he gets the chance he says: “*Exposure to tobacco smoke does not increase the risk of lung cancer and heart disease*”. Such statements are regularly qualified as bullshit⁷.

Such examples could easily be multiplied as there are a wide range of scientific theories that have pseudoscientific look-alikes. Of course, some producers and disseminators of pseudo-scientific content are real charlatans, and do not care

⁶ According to Hansson (2013, 70–71) a statement is pseudo-scientific if it satisfies the following three conditions (a) it belongs to a subject within the domain of science (b) it doesn’t respect the epistemic desiderata of sciences and lacks epistemic reliability, (c) it is part of a doctrine whose main proponents try to create the impression that it is scientific (i.e. the impression that it is the most reliable knowledge claim on the subject matter). The third condition explains why pseudo-scientific statements are so treacherous. Given that most lay people are not qualified to understand and assess scientific theories, but are looking for “cues that signify scientific competence by proxy” (Blancke, Boudry, and Pigliucci 2017: 87) pseudo-scientific claims dress up so to look reliable, often by mimicking superficial features of science like mathematical lingo and formulas, references to scientific theories and results, graphs and technical jargon.

⁷ For example, Italian virusologist Roberto Burioni is quoted in *Science* as calling bullshit many statements made by anti-vaccinists (Starr 2020). Likewise reporting on the Informed Health Choices Project, a project that aims to develop tools for improving critical health literacy in children, Matt Oxman describes pseudo-scientific claims as bullshit (Oxman 2016). Similar examples are gathered and discussed by Bergstrom and West (2020) and Yasmin (2021). For a discussion of pseudoscientific claims see Hansson (2013, 2017), Diethelm and McKee (2009), Ong and Glantz (2001), Oreskes and Conway (2010), Weigand (2021) and Wise (1998).

about the evidence for their claims. But then there are their honest followers, those duped and ensnared into believing and repeating their statements. These people are not indifferent toward the truth of their contribution, nor are many of them indifferent toward the evidence for their conversational contribution. Rather, they are so ignorant of the relevant subject that they don't even know what counts as adequate evidence in that field of knowledge. They fit Cohen's description of the *honest bullshitter* as the honest person who reads some bullshit believes it to be the truth, and later on affirms it (Cohen 2012: 332).

A *second example* that suggests a departure from analyzing bullshit in intentional terms is provided by Wreen:

“Imagine that after years of study I come up with a complicated system for beating the casinos that I sincerely believe is flawless. I travel across the United States lecturing about it to various groups, enthusiastically touting its virtues. In fact, I couldn't be more wrong: the system is seriously defective and contains multiple errors, silly even egregious errors.” (Wreen 2013: 110)

Wreen claims that whenever he utters sentences like “*My system will help you beat the casino*” his utterances are bullshit even though he doesn't believe he's bullshitting and certainly *doesn't intend* to bullshit. In this case, he is not aiming at fooling someone about the system, he is merely in *massive error*. I agree with his diagnostic, that this is a case of bullshit and I believe that many other examples can be found, especially if we look at utterances made by people who fall prey to quackery come to firmly believe that they hold the truth and want to enlighten everyone else⁸. Yet, Wreen's naïve inventor doesn't seem to be somebody who is indifferent toward the evidence for his claims, but just like the honest disseminators of pseudo-scientific content he is massively careless with the evidence for his claims.

Then, a proper characterization of bullshit is the following: If by utterance *u*, *S* communicates that *p* (where *p* is either what is said or the conversational implicature of *u*) *S* is bullshitting if and only if is careless toward the evidence for *p*.

As White (1961: 592–593) observes, there is a difference between being indifferent and being careless. A careless person is one who doesn't take care of something, while an indifferent person is one that doesn't care about that. Indifference is an attitude that one takes towards something. Its opposite is also an attitude: it is to care. For example, I might be indifferent if I pass a deadly virus to someone (i.e. I don't care). Or, on the contrary, I might care if I pass the virus to someone. But even if I care about this, I might be careless in my behavior: for

⁸ If one believes that a key feature of conspiracy theories is that they are epistemically faulty in that they are based on guess work rather than solid evidence, as Cassam (2019) argues, we can add disseminators of conspiracy theories to this list.

example, I might omit to take proper precautions or I might fail in other respects. In other words, indifference is an attitude that a person can take towards something, but carelessness is simply a failure to take care of things of which that person ought to take care. As such, indifference is intentional while carelessness is not. It follows, as White points out (1961: 593) that actions that are the result of indifference are to be explained in intentional terms while actions that are the result of carelessness are to be explained in terms of omissions.

Thus, there is a difference between speaking with an indifference toward the evidence for one's contribution and speaking with carelessness toward the evidence. Carelessness about the evidence can be the result of indifference toward the evidence, but it need not be so -it can also be the result of an omission. The converse is not valid though: if one is indifferent towards the evidence, then one is bound to be careless with the evidence. Hence this is an even more encompassing characterization of bullshit. The examples provided by utterances of honest disseminators of pseudoscientific content and Wreen's naïve inventor fit this characterization.

If we characterize bullshit in this more encompassing way, we can capture both the pedestrian bullshit analyzed by Frankfurt and make sense of the 'honest bullshit' that interested Cohen. Bullshit is speaking without giving due care to the evidence for what one is communicating. Some might perform this in full awareness of their lacking, and intentionally mislead their audience about their enterprise, just as many charlatans do. But others might perform it without being aware of their lacking, as is the case with the true-believers of pseudoscientific content and honest disciples of various charlatans. The former case counts as genuine deceit, while the latter need not. On this account, bullshit aligns with misleading rather than with deceiving when it comes to intention⁹.

We might wonder whether the honest believer, the mere disseminator of bullshit, is at any fault. The answer is 'yes', their fault is both epistemic and conversational and lies in being epistemically careless, in not paying attention to check the veracity of their claims. What counts as giving due care to evidence might vary with the context and the subject matter. If one recounts the bickering between two political parties, having read about it in a reliable newspaper counts as having adequate evidence, but if one passes judgment on the viability of a

⁹ Following Carson (2010: 42), Saul (2012: 72), Mahon (2016: 40) I take deception to be intentional, while misleading *can be* accidental (unintentional). Roughly, to mislead is to cause "false beliefs either intentionally or unintentionally" (Mahon 2016: 40) while to deceive is to "intentionally cause another person to acquire a false belief, or to continue to have a false belief, or to cease to have a true belief, or to be prevented from acquiring a true belief" (Mahon 2016, 44–5).

scientific hypothesis, what counts as having adequate evidence is to obey the norms of inquiry in the respective scientific field.

4 Quality: Truthfulness and reliability

The bullshitter's fault is not just epistemic but also conversational, and we can analyze it in Gricean terms as a violation of the second Quality maxim: *do not say that for which you lack adequate evidence*¹⁰. According to Grice, conversation is a form of joint-action and successful communication requires cooperation among interlocutors. Joint-actions bring about expectations and obligations on the part of the actors: we are *obliged* to act in certain ways if the joint action is to succeed, and we *expect* our partners to do their part for the action to succeed, and vice-versa. This is also true for conversations: in order for conversation to work, it requires many coordinating moves between interlocutors, but most importantly it requires them to be cooperative. In a sense, this is what the Cooperative Principle and the subordinate conversational maxims encapsulate: we expect our interlocutor to make her contribution truthful, relevant, to give the right amount of information required at that stage given the common goal of the conversation, and we oblige to do the same (Grice 1989: 26–27).

Among the maxims, the one relevant here is the supermaxim of Quality (*Try to make your contribution one that is true*) and the two more specific maxims, the first of which regards truthfulness (*Do not say what you believe to be false*) and the second which regards reliability (*Do not say that for which you lack adequate evidence*). If a speaker is cooperative she will try to make her contribution truthful and reliable. Obviously, Grice is not claiming that interlocutors are always genuinely cooperative, since time and again we are reminded that people use language to deceive and selfishly manipulate their audience. Rather, what he claims is that for a conversation to work, that is to achieve its commonly agreed goal, the speaker is expected to be cooperative.

Given that Grice reconstructs conversational exchanges as activities governed by norms, namely by the conversational maxims, he takes the maxims to impose certain normative requirements on the speaker. In a conversation, the audience is entitled to assume, in the absence of reasons for believing otherwise, that the

10 I understand a violation of a maxim to be its non-fulfilment that is not the result of opting out, nor of flouting nor of clashing with other maxims, and it can be intentional or unintentional. When one violates a maxim, one is basically non-cooperative, as opposed to when one flouts a maxim, or when what one says brings about a clash of maxims which are both compatible with being cooperative.

speaker's contributions correspond to the requirements set by the maxims. As Ross (1986:77–78) points out with respect to rule-governed behavior in general, “where the rules are such that one may perform a certain action only if a certain condition obtains, for example, that one may walk off with a suitcase only if it belongs to one or wear a plain gold ring on the third finger of the left hand only if one is married, then to perform the action is to entitle witnesses that the corresponding condition obtains”. Something similar occurs in conversations: they give rise to certain obligations, on the part of the speaker, and certain entitlements, on the part of the hearer. Again, the obligation that the supermaxim of Quality, together with the two submaxims, impose on the speaker is that her conversational contribution should be truthful and reliable. Conversely, the audience is entitled to trust the speaker to be sincere and reliable, and adopts, as a prerequisite for a fruitful exchange, a stance of trust in speakers' truthfulness and reliability, even if this trust is not blind and can easily be rescinded¹¹.

Put another way, when one is accepted in a norm governed activity one is trusted to exercise the responsibility that comes with taking part in that activity. As far as conversations are concerned, part of that responsibility is to have adequate evidence for the conversational contribution, given the subject matter, the goal and the standards of the conversation. For this reason, in regular conversations, contributions are expected to be reliable, and by being careless with the evidence the bullshitter undermines the reliability of her contributions and thus misleads (intentionally or not) her audience.

That interlocutors adopt a stance of trust in the speaker's truthfulness and reliability is also suggested jointly by data from hedging and discursive evidentiality¹². Data analyzed by Fetzer and Etsuko suggest that in languages where evidentiality is not grammatically encoded, ‘evidence’ occurs relatively rare in ordinary discourse and it is “generally brought into discourse when the conversational contribution is not accepted as given” (Fetzer and Otsuko 2014:327) or where there are chances that it will be contested. This is concurred by Sbisà's observation that “any doubt on the audience about the speaker's competence and overall reliability might lead to disbelief in [her contribution] and disqualification of the speaker as informant [...] such that her “contributions do not longer count as testimony and, therefore, the speaker faces a challenge” (Sbisà 2014: 472). One way for the speaker to prevent doubts from arising is to “boost her competence and

¹¹ As Sperber et al. (2010: 360–69) point out, audiences often use a certain amount of vigilance with respect to what they are being told. This vigilance is directed both towards the source and towards the content of what they are told.

¹² Evidentiality (as a dimension) can be described as expressing speaker's grounds for her conversational contribution.

reliability [by accompanying] her assertions with introductory or accessory comments aimed at displaying her credentials” (Sbisà 2014: 471). This can be done by various expressions, such as “according to *Encyclopedia Britannica*”, “According to the expert consensus on ...”, “I have seen ...”, “I have observed that ...”, etc. In the opposite direction, when the evidence one has is not fully adequate (either because its source is not completely reliable, or because it doesn’t fully support the contribution), one is expected to hedge their assertion by means of evidential hedges (Benton and van Elswyk 2020). By using evidential hedges, a speaker signals that the evidence for her contribution is less than fully reliable or is regarded to be as such. A diverse number of expressions can be employed: *it might be, perhaps, there is a chance, apparently, reportedly*, etc.

Taken together these data suggests that when unqualified sentences are uttered, the speaker’s reliability for her contribution is simply assumed. Or to put it in Gricean terms: interlocutors assume that the speaker conforms to the second maxim of Quality, that she has adequate evidence for her contribution - a fact that directly follows from the assumption that the speaker is cooperative. Exactly how we should understand the demand to have adequate evidence is discussed in the next section.

5 Alternative accounts and putative counterexamples

In this section I will discuss two accounts¹³ that bear certain resemblances with the one presented here, in order to set them apart, and to give further support to my view. I will also discuss and rebut two counterexamples raised against analyzing bullshit in Gricean terms.

5.1 What is adequate evidence?

Dynel (2011 and 2018) has also proposed that we characterize bullshit as a violation of the second maxim of Quality¹⁴. But her account is markedly different than the

¹³ There are many other theoretical accounts of bullshit that are worth discussing (Meibauer 2016, 2018, Wreen 2013, Saul 2012, Carson, 2016) but reasons of space prevent me from doing that here. I will leave this to further work.

¹⁴ Fallis (2009: 30–31) has also proposed analyzing bullshit as a violation of the second Quality maxim but since he has in meantime changed his mind, I will not discuss his position here. His new account is discussed in 5.2.

one presented here. The first difference is methodological: she advises us to understand bullshit as a technical term, stick with the Frankfurtian definition and then see how pervasive it is (given this definition) and how to capture it in pragmatic terms. In other words, she advises us against giving alternative definitions of bullshit because they merely “obfuscate the picture of bullshit as a category of deception” (Dynel 2018:335). It is no surprise then, that she believes that bullshit is less prevalent than Frankfurt and others have thought (Dynel 2018:329). I, on the other hand, have taken the opposite direction: to look at assertions that are generally taken to be bullshit, and subsume them under a unified account. More importantly, she takes bullshit to be essentially a form of deception and thus to necessarily involve the intention to deceive, while I do not. But the central point of difference lies in how we should interpret the second Quality maxim – *Do not say that for which you lack adequate evidence*. She proposes that we interpret “adequate evidence” as *the evidence that is adequate by the lights of the speaker*. “The adequacy of the evidence is decided on solely by the speaker” and “the evidence is [...] dependent on the speaker’s intentions and belief system rather than hard facts she says (Dynel 2018: 335). In contrast, I propose that what counts as adequate evidence is contextual but not subjective: it is determined by the subject matter and contextually salient epistemic standards, an approach in line with various forms of contextualism in epistemology (see Pynn 2017).

For reasons having to do with the demands of the Quality Supermaxim, namely sincerity and reliability, we should not to accept Dynel’s proposed interpretation of what counts as adequate evidence. If her subjectivist proposal is correct, there are two ways to interpret the second maxim:

- (a) “Do not say that for which you *believe* you lack adequate evidence”.
- (b) “Do not say that which does not conform to your standard of adequate evidence”

If we understand the second maxim as in (a), then the only way to violate it is to communicate a proposition for which you *believe* you lack adequate evidence. In that case, communicating a proposition for which objectively there is no adequate evidence but for which you believe you have adequate evidence is not a violation of the second maxim just as communicating something false which you believe to be true is not a violation of the first maxim of Quality. On this understanding, all violations of the second maxim of Quality must be *intentional*, and thus all amount to *insincerity*, just as violations of the first maxim of Quality amount to *insincerity*. But, on this understanding we lose an important aspect encoded in the supermaxim of Quality: *Try to make your contribution one that is true*. When we are required to *try* to make contributions that *are true*, we are required not just to make sincere contributions (this is not in the realm of trying), but we are also required to

try to make reliable contributions, since reliability increases the chances that a *truthful, or sincere*, contribution is also a *true* one. That is to say that speakers are expected to be not just sincere, but also reliable. When we are given a piece of information we expect the speaker to be both *reliable* (has adequate evidence for the information given) and *sincere* (she is willing to share that information with us, as opposed to intentionally mislead us by giving us information that she believes to be false and/or it believes to be ungrounded)¹⁵.

We should also resist the interpretation given in (b). Given that the two Quality maxims help put in practice the demands of the Quality Supermaxim, I believe there are good reasons why Grice has formulated them the way he did. The first Quality maxim demands that the speaker does not contribute propositions she *believes* to be false, and does *not* demand that the speaker does not contribute false propositions. After all, given that we're fallible, someone might communicate a false proposition and believe it to be true. Then, why not formulate the second Quality maxim in a subjective manner as well: do not say that which does not conform to your standard of adequate evidence? The reason has to do with the demands that the Cooperative Principle and its subordinate maxims place on the speaker. These demands are not subjective. The Principle demands that the speaker makes her "contribution as is required [...] by the accepted purpose or direction" of the conversation (Grice 1989, 26). And when, for example, some of these demands are spelled out by the two Quantity maxims they are to "make your contribution as informative as required (for the current purpose of the exchange) and "not to make your contribution more informative than required" (Grice 1989, 26). Notice that the demand is not put in subjective terms: make your contribution as informative as it is by your standards. Despite her good intentions, one might not fulfil the Quantity supermaxim, if one is not aware of the right amount of information that is required from her by the particular conversation she is engaged in.

For these reasons, I believe we should understand the condition of adequate evidence required by the second maxim to be contextually determined, and not subjective: the subject matter, the goal and the implicit epistemic standards determine how much evidence one needs in order to have adequate evidence. If this is so, a speaker can conversationally contribute a proposition for which she lacks adequate evidence (according to the evidential standards of the context) even though she believes she has.

¹⁵ Origgi (2004: 65) and Sperber et al. (2010: 369) argue that epistemic trust is two-folded: it includes expectation of truthfulness and expectation of reliability.

5.2 Bullshit and inquiry

Fallis and Stokke (2017) give an account of bullshit that is not coached in Gricean terms, but it bears some similarities to the one presented here, so it is worth discussing it.

Fallis (2014: 3) gives this characterization of bullshit: “an assertion is bullshit if and only if the speaker does not care whether the inquiry gets closer to the truth, gets further from the truth, or stays in the same place”. Building on earlier work by Stalnaker (1984) and Roberts (2012), Fallis and Stokke take communication to be a cooperative activity of information sharing aimed at inquiry: to discover how things *actually* are. The goal of conversation “is to incrementally contribute true information to a body of shared information, the *common ground*, with the aim of reaching a maximal information state that rules out all other possibilities than the actual world itself” (Fallis and Stokke 2017: 285). But most often, we are not engaged in an all-out inquiry into what is the world like, but rather into more immediate and local inquiries, such as finding out whether a vaccine is efficient against a particular virus, how cold it is outside, who is responsible for a terrorist attack, whether someone kept their end of the bargain, and so on. Such a sub-inquiry constitutes the topic of a particular conversation and it can be represented as a question that was explicitly or tacitly accepted as under discussion. Then on this view, conversations have “as immediate goals that of answering a set of questions under discussion” (Fallis and Stokke 2017: 286) – e.g. Is vaccination effective against MMR? Are summers, on average, rainy in Normandy? Did John keep his promise? etc.

Given this framework, they give the following definition of bullshit: “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” (Fallis and Stokke 2017, 295). In simpler words, someone is bullshitting if and only if, she is indifferent toward the evidence for her contribution and thus indifferent on whether her contribution takes the inquiry closer to the truth or moves it further from it.

This is a powerful and elegant view and it can account for a large number of cases, far beyond Frankfurt’s original examples. As I see it, it has one shortcoming, though: it cannot account for all types of honest bullshitting¹⁶. It cannot account for the bullshit produced by the disseminator of pseudo-scientific doctrines or of conspiracy theories, the one who honestly believes what she preaches and is genuinely preoccupied with finding how things are. Nor can it

¹⁶ Fallis and Stokke explicitly admit this (2017: 295).

account for the overconfident person who grounds his statements on guts alone. It is quite sure that some anti-vaccinists intend their statements, (for example that “Vaccines are not safe and cause autism”) to contribute to the inquiry, as some disseminators of conspiracy theories do, and such statements are regarded to be bullshit (see Bergstrom and West (2020) and Yasmin (2021) and references in footnote 9). But by Fallis and Stokke’s lights, such statements do not count as bullshit. On the other hand, Carson’s example of evasive speech does come out as bullshit in their account, since the speaker is not interested in contributing to the goal of inquiry while on my account, in agreement with Meibauer (2016), I take evasiveness to be an altogether different category of misleading speech. I don’t believe there is a knockdown argument in favor of one or another position, but Fallis and Stokke need to give an alternative account for the examples mentioned above, and they need to explain why regular folk err in regarding those as examples of bullshitting.

5.3 Evasive bullshit

An example that allegedly counts against treating bullshit as carelessness towards evidence is put forward by Carson (2010: 61–62). He discusses a situation where one says something that he knows to be false, and therefore has good evidence of its falsity, but which, claims Carson, is a form of bullshit.

Suppose that I teach at a university that is very intolerant of atheists. I am asked by an administrator whether a friend and colleague is an atheist. I know that he is an atheist and that it will harm him if I reveal this. I do not want to harm my friend nor do I want to lie and say that he is not an atheist as I fear that I am likely to be found out if I lie about this. I give an evasive bullshit answer. I say “as a boy he always went to Church and loved singing Christmas Carols” even though I know this to be false. (I am not worried that I will be caught or found out if I lie about this) (Carson 2010: 61–2).

I would contest that this is a case of bullshit, although it is an evasive and misleading answer. In this case the speaker lies and misleads by untruthful implicature at the same time. He says something that he believes to be false (i.e. that as a boy his friend attended Church and loved singing Christmas Carols) with the intention to make his audience believe that is true – hence the lying. He also exploits the maxim of Relation by giving an irrelevant answer to the administrator’s question, thus inviting the false implicature that his friend is not an atheist. This is not a rare case. It is not uncommon that speakers are evasive and evasiveness involves the exploitation of the Relation Maxim (maybe in conjunction with other maxims) either to invite

untruthful conversational implicatures or simply to convey that the speaker refuses to engage with what was previously said¹⁷.

5.4 Lackey assertions

Stokke (2018a) offers another example that, he claims, speaks against modeling bullshit in neo-gricean terms. He borrows Lackey's "Creationist Teacher" case:

Stella is a devoutly Christian fourth-grade teacher, and her religious beliefs are grounded in a deep faith that she has had since she was a very young child. Part of this faith includes a belief in the truth of creationism and, accordingly, a belief in the falsity of evolutionary theory. Despite this, she fully recognizes that there is an overwhelming amount of scientific evidence against both of these beliefs. Indeed, she readily admits that she is not basing her own commitment to creationism on evidence at all but, rather, on the personal faith that she has in an all-powerful Creator (Lackey 2008: 48).

Stokke (2018a: 229) asks us to suppose that Stella is asked, in a private conversation, outside school, what she thinks about the origin of species, to which she replies, "God created the species." Stokke claims that "despite the fact that Stella says and intends to communicate something for which she believes she lacks adequate evidence; she is surely not bullshitting".

But there are reasons to resist Stokke's argument. One consideration that goes against his argument is the following. Intuitions on this particular case are muddled with what one tacitly believes about the affirmation of propositional faith. More precisely, how one judges this particular case depends on their tacit attitude towards the affirmation of propositional faith: is it an assertion or some other form of non-assertoric speech act? To see this, consider another context: suppose Stella takes part in a theoretical biology seminar and she is asked what she thinks about the origin of species, to which she replies "God created the species". Would we count what she said to be bullshit? Most likely, yes. Now, consider a third context: as a teacher, Stella is invited to a trial to give her opinion on introducing the so-called intelligent design theory into the curricula of biology classes. When asked what she thinks about the origin of species she replies "God created the species". Would we count what she said to be bullshit? Again, most likely, yes. Why would things be different in the context of a private conversation,

¹⁷ As Meibauer points out (2018, 365) this seems to be no different than Grice's famous example of evasive speech where A says "Mrs. X is an old bag." And B replies "The weather has been quite delightful this summer" inviting the implicature that A's remark should not be discussed (Grice 1989: 35). Meibauer (2016; 2018) argues correctly, to my mind, that evasive bullshit is just a particular case of evasiveness that results from flouting the Maxim of Relation.

where she merely wants to inform the audience about her religious beliefs? An explanation is that in the seminar and the trial contexts we take Stella to make a *knowledge claim*, thus a claim which is subject to evidential requirements. In the private context we do not take her to make a knowledge claim, but merely to *affirm her faith*. Those like Stokke who do not consider Stella's utterance in the private context to be bullshit are tacitly committed either to a form of hermeneutic fictionalism about religious discourse, or to a form of non-evidentialism about religious belief. Some understand religious affirmations as non-assertoric speech acts that only superficially resemble assertions but do not commit speakers to the truth of what is said – maybe they are confessional acts rather than expressions of belief, or a form of non-doxastic praise – while others are not evidentialists about religious beliefs, and do not take an affirmation of religious belief to be governed by evidential norms¹⁸ (see Scott 2020). In any case, I do not intend to suggest that we can draw a clear line between contexts where we count an utterance of “God created the species” as a mere affirmation of faith and contexts where it is subject to evidential requirements. But whether a statement like Stella's is regarded as bullshit depends on whether one views her statements to be affirmation of faith or not, and on further views on religious discourse and the epistemology of religious belief.

6 Concluding remarks

In this paper I have given a characterization of bullshit that is more encompassing than Frankfurt's original characterization. In a sense, the one given here

18 Here I depart from Lackey's original example, where Stella, the creationist teacher, regards her duty to present only material that is best supported by the available evidence, which includes the truth of evolutionary theory, and thus never states in class that God created the species. But Kurt Wise is a real life example: a science professor and young-earth Creationist who by his own admission came to admit young-earth theory not because of scientific evidence, but in spite of it: “although there are scientific reasons for accepting a young Earth, I am a young-age creationist because that is my understanding of the Scripture. As I shared with my professors years ago when I was in college, *if all the evidence in the universe turns against creationism, I would be the first to admit it, but I would still be a creationist* because that is what the Word of God seems to indicate. Here I must stand.” Either the Scripture was true and evolution was wrong or evolution was true and I must toss out the Bible ... It was there that night that I accepted the Word of God and rejected all that would ever counter it, including evolution. With that, in great sorrow, I tossed into the fire all my dreams and hopes in science.” (quoted in Wise 1998 - emphasis added). When Kurt Wise says, in class or in public interviews “God created the species”, *pace* Stokke, we regard his claims as bullshit.

can be seen as an extension of Frankfurt's, since his definition comes out as a particular case of the one defended here. But we should depart from Frankfurt's definition in two respects: (1) we should not characterize bullshit in relation to truth; and (2) we should not characterize it in intentional terms. Rather, bullshit is to be characterized as speaking with carelessness toward the evidence for one's conversational contribution, in conversational situations where one is expected to be reliable. I have argued that this characterization can be modeled in Gricean terms as a violation of the second maxim of Quality: *Do not say that for which you lack adequate evidence* and that we should understand this maxim as demanding on the speaker that her conversational contribution is epistemically reliable. Which is to say that in a conversation the audience assumes that the speaker is observing the demands of the Supermaxim of Quality and thus assumes that the speaker is sincere (believes what she is saying to be true) and reliable (has adequate evidence for her conversational contribution), where what counts as adequate evidence should not be understood in subjective terms (i.e. whatever is adequate by speaker's standards) but in contextual terms: the subject matter and the context's epistemic standards determine what counts as adequate evidence. The examples discussed in Section 3 suggest that the analysis proposed here is supported by our common usage of "bullshit" as a label for speech performed with carelessness toward the (contextually adequate) evidence for one's contribution. Basically, the bullshitter fails to be an epistemically reliable interlocutor and betrays the audience's trust and expectation of reliability. Among other things, this approach also cashes out Frankfurt's observation that "bullshit is unavoidable whenever circumstances require someone to talk without knowing what he is talking about" (Frankfurt 2005, 63). If in Dynel's (2018) analysis bullshit comes out as less prevalent than Frankfurt believed, according to my analysis just the opposite might be true. A disquieting consequence of this view might be that "we all bring our contribution to bullshit", as Frankfurt observes at the very beginning of his essay. This is so, I venture to say, because we are all, from time to time, careless with the evidence for our conversational contributions.

Acknowledgments: I would like to thank the two anonymous reviewers of this journal as well as Tobies Grimaltos, Andrei Moldovan, Bianca Savu, Sergiu Spătan and Dan Zeman for their helpful comments and suggestions on previous drafts of the paper. I also thank audiences at Complutense University of Madrid, University of Valencia and University of Bucharest for their questions, comments, objections and suggestions.

References

- Blanche, Stefaan, Maarten Boudry & Massimo Pigliucci. 2017. Why do irrational beliefs mimic science? The cultural evolution of pseudoscience. *Theoria* 83(1). 78–97.
- Benton, Matthew A. & Peter Van Elswyk. 2020. Hedged assertion. In Sanford Goldberg (ed.), *The oxford handbook of assertion*, 245–263. Oxford: Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780190675233.013.11>.
- Bergstrom, Carl & Jevin West. 2020. *Calling bullshit*. New York: Random House.
- Carson, Thomas L. 2010. *Lying and deception*. Oxford: Oxford University Press.
- Carson, Thomas L. 2016. Frankfurt and Cohen on bullshit, bullshitting, deception, lying, and concern with the truth of what one says. *Pragmatics and Cognition* 23(1). 53–67.
- Cassam, Quassim. 2019. *Conspiracy theories*. Cambridge: Polity Press.
- Cohen, G.A. 2012. Deeper into bullshit. In G.A. Cohen (ed.), (2012) *Finding oneself in the other*. Princeton NJ: Princeton University Press.
- Davis, Wayne A. 2021. Cognitive propositions and semantic expressions. *Intercultural Pragmatics* 18(3). 337–358.
- Diethelm, Pascal & Martin McKee. 2009. Denialism: What is it and how should scientists respond? *The European Journal of Public Health* 19(1). 2–4.
- Dynel, Marta. 2011. A web of deceit: A neo-gricean view on types of verbal deception. *International Review of Pragmatics* 3. 139–167.
- Dynel, Marta. 2018. *Irony, deception and humour: Seeking the truth about overt and covert untruthfulness*. Berlin, Boston: De Gruyter Mouton.
- Fallis, Don. 2009. What is lying. *Journal of Philosophy* 106(1). 29–56.
- Fallis, Don. 2014. Frankfurt wasn't bullshitting! *Southwest Philosophical Studies*. Available at: <https://ssrn.com/abstract=2462769>.
- Fallis, Don & Andreas Stokke. 2017. Bullshitting, lying, and indifference toward truth. *Ergo: An Open Access Journal of Philosophy* 4(10). <https://doi.org/10.3998/ergo.12405314.0004.010>.
- Fetzer, Anita & Etsuko Oishi. 2014. Evidentiality in discourse. *Intercultural Pragmatics* 11. 321–332.
- Frankfurt, Harry G. 2002. Reply to Cohen. In Sarah Buss & Overton Lee (eds.), *Contours of agency: Essays on themes from Harry Frankfurt*. Cambridge Mass: MIT Press, Bradford Books.
- Frankfurt, Harry G. 2005. *On bullshit*. Princeton, NJ: Princeton University Press.
- Grice, H. Paul. 1989. Logic and conversation. In H.P. Grice (ed.), (1989) *Studies in the way of words*. Cambridge Mass: Harvard University Press.
- Hansson, Sven Ove. 2013. Defining pseudo-science and science. In Massimo Pigliucci & Maarten Boudry (eds.), *Philosophy of pseudoscience: Reconsidering the demarcation problem*. Chicago: University of Chicago Press.
- Hansson, Sven Ove. 2017. Science denial as a form of pseudoscience. *Studies in History and Philosophy of Science Part A* 63. 39–47.
- Lackey, Jennifer. 2008. *Learning from words: Testimony as a source of knowledge*. Oxford: Oxford University Press.
- Mahon, James Edwin. 2016. The definition of lying and deception. In Edward N. Zalta (ed.), *The stanford encyclopedia of philosophy*, Winter 2016 Edition. Available at <https://plato.stanford.edu/archives/win2016/entries/lying-definition/>.

- Meibauer, Jörg. 2016. Aspects of a theory of bullshit. *Pragmatics and Cognition* 23(1). 68–91.
- Meibauer, Jörg. 2018. The linguistics of lying. *Annual Review of Linguistics* 4. 357–375.
- Ong, Elisa K. & Stanton A. Glantz. 2001. Constructing “sound science” and “good epidemiology”: Tobacco, lawyers, and public relations firms. *American Journal of Public Health* 91. 1749–1757.
- Oreskes, Naomi & Erik M Conway. 2010. *Merchants of doubt*. London: Bloomsbury Press.
- Origgì, Gloria. 2004. Is trust an epistemological notion? *Episteme* 1(1). 61–72.
- Oxman, Matt. 2016. *Immunizing children against deadly bullshit* In *Healthnewsreview.org*, 16 July 2016 <https://www.healthnewsreview.org/2016/10/informed-health-choices/#.X1D1zeerF4c.link> (Accessed 9 March 2020).
- Pynn, Geoff. 2017. The intuitive basis for contextualism. In Jonathan Jenkins Ichikawa (ed.), *The routledge handbook of epistemic contextualism*, 32–43. Routledge. <https://doi.org/10.4324/9781315745275-3>.
- Roberts, Craige. 2012. Information structure in discourse: Toward an integrated formal theory of pragmatics. *Semantics & Pragmatics* 5. 1–69.
- Ross, Angus. 1986. Why do we believe what we are told? *Ratio* 1. 69–88.
- Rucker, Philipp & Dawsey, Josh and Daletta, Damian. 2018. Interview with President Trump. *Washington Post*, November 28, 2018.
- Saul, Jennifer. 2012. *Lying, misleading, and what is said: An exploration in philosophy of language and in ethics*. Oxford: Oxford University Press.
- Sbisà, Marina. 2014. Evidentiality and illocution. *Intercultural Pragmatics* 11(3). 463–483.
- Scott, Michael. 2020. Faith, fictionalism and bullshit. *Thought: A Journal of Philosophy* 9(2). 94–104.
- Sperber, Dan, Fabrice Clément, Christophe Heintz, Olivier Mascaro, Hugo Mercier, Gloria Origgì & Deirdre Wilson. 2010. Epistemic vigilance. *Mind & Language* 25(4). 359–393.
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge Mass: MIT Press.
- Starr, Douglas. 2020. Fighting words. *Science* 367(6473). 16–19.
- Stokke, Andreas. 2018a. Bullshitting. In Jörg Meibauer (ed.), *The oxford handbook of lying*. Oxford: Oxford University Press.
- Stokke, Andreas. 2018b. *Lying and insincerity*. Oxford, UK: Oxford University Press.
- Webber, Jonathan. 2013. Liar! *Analysis* 73(4). 651–659.
- Weigand, Edda. 2021. Language and dialogue in philosophy and science. *Intercultural Pragmatics* 18(4). 533–561.
- White, Alan R. 1961. Carelessness, indifference and recklessness. *The Modern Law Review* 24. 592–595.
- Wise, Donald. 1998. Creationism’s geologic time scale: Should the scientific community continue to fight rear-guard skirmishes with creationists, or insist that “young-earthers” defend their model in toto? *American Scientist* 86(2). 160–173.
- Wreen, Michael. 2013. A P.S. on B.S: Some remarks on humbug and bullshit. *Metaphilosophy* 44(1–2). 105–115.
- Yasmin, Seema. 2013. *Viral BS. Medical myths and why we fall for them*. Baltimore: John Hopkins University Press.

Bionote

Adrian Briciu

Department of Philosophy and Communication Sciences, West University of Timisoara, Timisoara, Romania

adrian.briciu@gmail.com

Adrian Briciu (PhD University of Barcelona, Spain) is assistant professor at West University of Timisoara, Romania. His research interests focus on indexicality, contextual effects on compositionality, the semantics-pragmatics interface and, lately, on the pragmatics of misleading and deceitful speech. He has published on the semantics and pragmatics of indexical expressions, semantic underdetermination and pragmatic enrichment and on the semantics of context-shifting expressions.