

## Bullshit Assertion

Kotzee, Ben

*License:*

None: All rights reserved

*Document Version*

Peer reviewed version

*Citation for published version (Harvard):*

Kotzee, B 2019, Bullshit Assertion. in SG (ed.), *Oxford Handbook of Assertion*. Oxford University Press.

[Link to publication on Research at Birmingham portal](#)

### General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

### Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact [UBIRA@lists.bham.ac.uk](mailto:UBIRA@lists.bham.ac.uk) providing details and we will remove access to the work immediately and investigate.

# **Bullshit Assertion**

Ben Kotzee, University of Birmingham

## **Introduction**

In his short article ‘On Bullshit’ (1986, republished as a book in 2005), Harry Frankfurt diagnoses a distinctive problem of contemporary culture: that so much of it is bullshit. Bullshit abounds in advertising, politics, the media and the academy; in fact, much communicative activity in our culture seems to be bullshit by design. But Frankfurt was worried about the volume of bullshit produced in our culture in particular; he blamed the sheer mass of bullshit that we encounter in public discourse in our society on the mass media, on advertising, on the party political system and on some currents in academic thought (notably post-modernism). More than diagnose a widespread problem however, in ‘On Bullshit’ Frankfurt developed a theoretical account of the nature of bullshit. In his account, Frankfurt holds that bullshit is, like lying, a dishonest assertion; however, he makes clear how bullshitting is different from lying and holds (startlingly) that bullshit is more dishonest than lying. In this chapter I consider bullshit as a distinct corruption or perversion of assertion next to lying. I hold that understanding bullshit and lying as perversions of assertion sheds light on assertion and how it functions.

## **Frankfurt’s account of bullshit**

Frankfurt credits Max Black with the first analysis of bullshit - what Black (1985) calls ‘humbug’.<sup>1</sup> Black defined humbug as:

deceptive misrepresentation, short of lying, especially by pretentious word or deed, of somebody’s own thoughts, feelings, or attitudes. (Black, 1983: 145)

Black describes humbug as a kind of misrepresentation. In making clear what kind, he draws a distinction between two matters that a speaker may mislead a listener about: the speaker may misrepresent what is the case in the world, or they may misrepresent their own feelings and attitudes to what is the case in the world. Black holds that the primary intention in humbug is not to deceive the listener about what is the case in the world but to deceive the listener about their thoughts, feelings or attitudes to it. Frankfurt builds his account of bullshit on Black’s view that bullshit is a kind of misrepresentation. Frankfurt notes that sometimes the misrepresentation involved in bullshit lies in communicating poorly thought out ideas in a hasty and sloppy fashion. Take the following example: the student did not study for an exam and the next day writes pages of whatever they can remember was said in class, using every item of remembered or misremembered technical vocabulary. This is called ‘bullshitting an exam’. One can also bullshit a class or a job interview or a sales presentation – in most of these cases, bullshit involves scanty preparation or sloppy communication laid over with confident presentation.

---

<sup>1</sup> Black also lists a number of synonyms for humbug: ‘balderdash’, ‘claptrap’, ‘hokum’, ‘drivel’, ‘buncombe’, ‘imposture’ and ‘quackery’. All of these can be subsumed under Frankfurt’s more direct ‘bullshit’.

Frankfurt points out, though, that bullshit is not always mere careless talk. Sometimes bullshit is very carefully crafted: he notes that in advertising, public relations and, in particular, in politics, people sometimes take great care to ensure that they will mislead others about their own feelings and attitudes without quite lying. Frankfurt asks us to

Consider a Fourth of July orator who goes on bombastically about “our great and blessed country, whose Founding Fathers under divine guidance created a new beginning for mankind.” (Frankfurt 2005: 16)

As Frankfurt notes, the Fourth of July orator most likely spent a great deal of time polishing and rehearsing their bombastic speech. This is in contrast to the student bullshitting the exam who did not prepare and just ‘throws words at the page’. Frankfurt acknowledges that there is a tension between bullshit as the covering up of shoddy or careless communication and bullshit as highly polished deception. (2005: 22) He holds, though, that there is always something substandard even to carefully crafted bullshit, because the bullshit artist is trying to ‘get away with something’ (or ‘faking it’) in some respect. While the fourth of July orator may take care over the style of their speech, she still does not take care about the *truth* of what she says.

The concern with truth forms the central plank of Frankfurt’s argument. It captures his disagreement with Black and is also the reason why we should take bullshit seriously in the context of understanding assertion. Firstly, Frankfurt distinguishes bullshit from lying. More fundamentally, in distinguishing bullshitting from lying, Frankfurt also distinguishes it from honest assertion. The most basic accounts of assertion explain asserting in terms of its social function of conveying of truth: what people do in making assertions is provide each other with information cooperatively. Because true belief is (usually!) instrumentally valuable, sharing a truth with another is to do the hearer the favour of passing useful information to them; conversely, asserting what is false is to do them the disservice of passing on bad information. One may assert what is false for two main reasons: one may be mistaken in what one believes, or one may deliberately say what one knows is false. Deliberately saying what one knows to be false is lying. (Fallis, 2009)

To lie to someone, like breaking a promise, is morally wrong and Kant famously held that we have an absolute duty not to lie.<sup>2</sup> However, in the philosophy of language and in epistemology, lying also receives considerable attention as a form of dishonest assertion or testimony.<sup>3</sup> A first point to note about lying is the extent to which it is dependent on, or parasitic on honest assertion. To see this, notice that the efficacy of lying depends on most people speaking honestly most of the time. In a world in which everyone always lied, it would be inadvisable to listen to what anyone has to say... and, if no-one listened, there would be little point speaking (telling lies) either. The point is that widespread lying would lead to the breakdown of communication and that the sometimes efficacy of lying actually depends on most people speaking honestly most of the time. This is to give a social argument – an argument from the practicality conditions of communication – regarding what is right about proper assertion and wrong about lying.

Drawing on arguments like these, a number of philosophers have offered accounts of the nature of assertion according to which assertion is essentially governed by a norm of

---

<sup>2</sup> For a discussion of the ethics of lying, see Bok (1978)

<sup>3</sup> In the field of assertion, there is a substantial literature on lying. See Adler (1997) and Carson (2006). In this volume, see T. Cuneo, “The Ethics of Assertion,” and J. Pepp, “Lying, Misleading, and Falsely Implicating”].

propriety to the effect that one should only say what is true, that one should only say what one believes, that one should only say what one knows, etc. and that what one should *not* do in asserting is lie. A number of possible norms of assertion have been proposed: that one should only assert what one believes, that one should only assert what is true, that one should only assert what is justified or that one should only assert what one knows and considerable debate exists in the literature regarding exactly what form the norm of assertion takes.<sup>4</sup>

While he does not provide an account of assertion himself, Frankfurt situates his account of bullshit in such normative thinking regarding assertion. In lying, Frankfurt holds, the liar is engaged in the *practice* of asserting just as the speaker who is speaking the truth is. By this, he means that the liar is transmitting to the hearer a discreet piece of information that is of interest to the hearer and that is portrayed as true. Importantly, in saying what he says the liar is *guided* by the truth. The liar could not lie if he did not have definite opinions about what is true and false and consciously avoided telling the truth: he is guided in what he says *by* what he takes to be true in as much as he tries to *avoid* saying what is true. (Frankfurt, 2005: 56 – 61)

In bullshitting, the situation is different. The bullshitter does not assert something that he believes is false. In reality, the bullshitter does not care much about what he actually says – his aim is not to lead the listener into believing something specific that is not true, but to confuse the listener into believing that he is communicating something *at all* when he is in fact communicating nothing specific. On Frankfurt's view, whereas the liar *informs* the listener (but wrongly), the bullshitter is faking taking part in the activity of information transmission. As Frankfurt says of the Fourth of July orator:

Her statement is grounded neither in a belief that it is true nor, as a lie must be, in a belief that it is not true. It is just this lack of connection to a concern with truth – this indifference to how things really are – that I regard as of the essence of bullshit. (Frankfurt 2005: 33 – 4)

For Frankfurt, the bullshitter's indifference to truth makes them 'a greater enemy of the truth than [the liar].' (2005: 61)

### **Criticism of the Frankfurt view**

There are three sorts of problems with Frankfurt's view that bullshitting consists in not caring about the truth: (1) there can be cases in which someone clearly cares about the truth and still bullshits, (2) there can be cases in which someone lies and bullshits at once and (3) there can be cases in which people seem to believe their own bullshit.

A number of authors<sup>5</sup> point out that sometimes the bullshitter does care about the truth or falsity of what they say. Take the case of the student asked an exam question that they do not know how to answer. Not knowing how to answer, they studiously avoid answering the exact question, but write a carefully crafted answer on some adjacent matters, trying to give the impression that they are answering the question when they are merely skirting around it. According to Carson (2010), the careful exam taker actually cares about the truth very much, rather than very little. This shows up in the fact that the student takes great care to write

---

<sup>4</sup> More on this matter below. For discussion, see the entry by M. Simion and C. Kelp in this handbook.

<sup>5</sup> E.g. Cohen, 2002; Kimbrough, 2006; Carson, 2010; Wreen, 2013; Stokke and Fallis, 2017.

something that is true (even though it does not answer the exact question) and avoids writing what is clearly false.

Closely related to the tactic of the careful exam taker is the tactic of misdirection. Jonathan Webber (2013: 652) cites as an example a famous *BBC Newsnight* interview between interviewer Jeremy Paxman and British Home Secretary Michael Howard. Paxman asked Howard whether he attempted to force Derek Lewis, Director of Her Majesty's Prison Service, to fire John Marriot, governor of the Parkhurst prison. In the interview, Paxman repeatedly asked Howard the exact same question: 'did you threaten to over-rule him?' and each time, Howard answered with a precise and correct but evasive answer:

Paxman: 'Did you threaten to over-rule him?'

Howard: 'I was entitled to express my views, I was entitled to consult him'

Paxman: 'Did you threaten to over-rule him?'

Howard: 'The truth of the matter is, Mr Marriot was not suspended'

Paxman: 'Did you threaten to over-rule him?'

Howard: 'I did not over-rule Derek Lewis'

In the end, Paxman asked Howard the same question 12 times and each time Howard responded by confidently sidestepping the question. Howard went so far as to admit that he 'gave him [Lewis] the benefit of my opinion in strong language', but pointedly refused Paxman's request to answer the question whether he 'threatened to overrule him' with a 'yes' or 'no'. Throughout the exchange, Howard took great care to say only what is strictly true and to avoid being caught out lying on air, even though it is obvious from the context that he was being knowingly evasive.<sup>6</sup> Webber (2013: 652) calls what Howard did 'filibustering' and holds that what Howard attempted to do was to *imply* that he acted properly, without actually *asserting* something that is false. By appeal to this example, Webber argues that one can imply something that is bullshit even while asserting something that is strictly true – contradicting Frankfurt's account that the bullshitter asserts something the truth of which they do not care about. Together, these two kinds of example (the careful exam taker and the filibustering politician) show that the bullshitter is not always indifferent to the truth – sometimes the bullshitter cares about the strict truth of what they say very much and may even bullshit by speaking the truth.

Moving on from Webber and Carson's accounts, if it is possible to bullshit by speaking the truth, it may even be possible to bullshit by lying. In fact, it is not hard to construct a case in which a person lies and bullshits at the same time or bullshits *by* lying. Take the following examples. To appear sophisticated in front of my hosts at the art gallery I opine:

'Yes, this installation is like the Eskimos who have 100 words for snow, you know: how it rejects a discredited aesthetic and establishes itself as beautiful in its own terms.'

Sounding off in a huddle of soldiers I say:

---

<sup>6</sup> Readers may watch the interview themselves on social media, e.g. here <https://www.youtube.com/watch?v=pyqnu6ywhR4> Howard held that the relevant matter was whether he was entitled or not entitled to instruct Lewis to fire Marriot. However, whether he *was* so entitled or not does not settle the question whether he actually threatened Lewis. At the end of the interview, Paxman rightly points out that Howard never answered the question.

‘It took a Mussolini to make the trains run on time. We must take no prisoners and come in for the big win.’

Neither of these things is true<sup>7</sup>, *and I know it*. However, both of these things are often said and I might choose to *lie* to this effect not because I particularly want the hearers to believe my lie, but because I want to imply something about myself (that I am culturally knowing and right-on in the first case, or that I am go-getting and tough in the second).

Third, we often say of people that they *believe* their own bullshit: by this, we mean that they have become so caught up with the grand but empty things that they say that they repeat it whenever opportunity presents. They may even repeat the bullshit utterances that they (or others) have made in the past seriously and may believe of themselves that this signals something clever and important about them as people. The problem for Frankfurt’s account is that, if bullshit involves a person *not* being serious about the truth or otherwise of what they say, it seems difficult to explain how someone may *believe* his/her own bullshit. After all, believing exactly does mean taking a proposition seriously as a personal commitment, standing by it, defending it, etc. The fact that people sometimes come to believe their own bullshit therefore presents a problem for Frankfurt.

### **Alternative accounts of bullshit**

Responding to the sorts of criticism outlined, a number of authors have offered alternative accounts of bullshit.

#### *Bullshit as nonsense*

Rather than offering an account that characterizes bullshit in terms of the speaker’s intentions when producing bullshit, Cohen offers an ‘output’- or ‘product’-based conception of bullshit. (Cohen, 2006: 121; 134)<sup>8</sup> Cohen holds that not all bullshit is produced with a bullshitting intention: some of it is bullshit by virtue of features of the utterance itself, independently of the speaker’s intention. Cohen distinguishes between the activity or process of bullshitting and the output of bullshitting – the shit that is uttered by someone who talks bullshit. Cohen is concerned with the bullshit that pervades academia; as an example he mentions the writings of Althusser and his followers. (2006: 118) He describes a situation that many academics are all too familiar with: reading pages of impenetrable jargon-laden text by some writer known for the complexity of his work, one sometimes begins to believe that there must be a deep point in the text that one is simply too dull to grasp. What interests Cohen is how it is possible that we are sometimes taken in by bullshit or why writing that is over-wrought and ambiguous in academia is often taken to be deep and complex. Writing about the Althusserian Marxists, Cohen says:

I found that material hard to understand, and, because I was naïve enough to believe that writings that were attracting a great deal of respectful, and even reverent attention could not be loaded with bullshit, I was inclined to put the blame for finding the Althusserians hard entirely on myself. And when I managed to extract what seemed like a reasonable idea from one of their texts, I attributed to it more interest or more

---

<sup>7</sup> There is widespread disagreement in linguistics and anthropology whether there are particularly many Eskimo words for ‘snow’. See Pullum, 1991. The punctuality of Italian trains did not improve during the time of Mussolini. See Dudley, 2016.

<sup>8</sup> In the recent literature, Ivankovic (2016) and Wreen (2013) offer accounts similar to that of Cohen.

importance... that it had, partly, no doubt, because I did not want to think that I had been wasting my time. (2006: 118)

Cohen holds that Frankfurt's analysis does not work for academic bullshit and that there are two things wrong with it: Firstly, entirely honest people often repeat on trust what other seemingly trustworthy people say; if these people talk bullshit, this bullshit is liable to be repeated, but, on Frankfurt's view, according to which bullshit requires dishonesty, as soon as the honest person repeats it, the utterance will stop being bullshit. Secondly, it is possible that an honest person may simply have a shitty idea and voice it; especially if this bullshit is repeated by others this person may even be encouraged to spout more similar sounding bullshit and others may begin to produce such bullshit of their own under the spell of the original. Because of this possible disconnect between the shittiness of what is said and the utterer's state of mind, Cohen holds that dishonesty is neither sufficient nor necessary for bullshit. In the end, Cohen holds that bullshit is a species of *nonsense*; he defines bullshit as 'unclarifiable unclarity'<sup>9</sup> (2006: 129) and hold that, whoever speaks nonsense – intentionally or not – is guilty of bullshitting.

### *Bullshitting by implication*

Above, we have already encountered Webber's (2013) account of bullshitting as misleading by implicature. Like Frankfurt, Webber draws a distinction between lying and bullshitting and provides an account of exactly what kind of misleading is involved in lying and in bullshitting. However, Webber is critical of Frankfurt's view that the 'bullshitter is a greater enemy of the truth than the liar'. Webber notes that, if Frankfurt is right, we should reserve greater opprobrium for the bullshitter who is caught out than the liar who is caught out. In point of fact, however, we tend to censure liars more heavily than bullshitters. This observation drawn from our actual linguistic practice does not square with Frankfurt's view that the bullshitter is a greater enemy of the truth than the liar.

Webber (2013: 655) acknowledges the existence of the kind of bullshit Frankfurt identifies; assertion that is indifferent to the truth of what is asserted; he calls this 'bullshit assertion'. However, Webber shows that there is another kind of bullshit, what he calls 'bullshit conversational implicature' that functions differently from bullshit assertion. For instance, above we encountered Webber's example of Michael Howard attempting to bullshit his way through the *Newsnight* interview. When the interviewer asks him 'did you threaten to overrule Mr Lewis', Howard says, truly, 'I did not overrule him'. While he only says that he did not *overrule* the Director of HM Prisons, he means to imply that he did not even threaten to overrule him either. The bullshit is on the level of the implicature, not on the level of what is actually said.

While Frankfurt holds that the bullshitter is a greater enemy of the truth than the liar, Webber holds that lying is more serious than bullshitting. Webber's argument has to do with the consequences of lying and the consequences of bullshitting. If a person is found out to lie, they thereby show themselves not to be trustworthy about what they say directly. However, if someone is found out to be bullshitting – like Howard was by Paxman's insistent questioning – their credibility is damaged in a different way: they are found out to be reliable on the strict level of what they *say*, but not on the level of what they *imply*. As Webber puts it:

---

<sup>9</sup> In similar vein Wren (2013) defines bullshit as 'noxious nonsense'.

... lying damages both credibility in assertion and credibility in implicature, whereas misleading damages only credibility in implicature. (2013: 651)

The point for Webber is that

We can still rely on [bullshitters], though we need to be careful, but cannot rely on liars. (2013: 654)

Webber holds that this distinction in terms of credibility damage is the reason why we take lying as a more serious epistemic offence than bullshitting. (2013: 658)

### *Bullshit answers to questions*

In an account that aims to extend the Frankfurt account, Stokke and Fallis (2017) focus on the malign role that bullshitting plays in inquiry. They hold that Frankfurt is right about indifference to truth being the essence of bullshit. However, they hold that there is more to indifference to truth than simply not caring whether a particular state of affairs  $p$  obtains or not. Stokke and Fallis stress that assertions do not tend to occur in isolation, but form part of discourse. They draw on Stalnaker's account of inquiry, according to which the aim of all discourse is to contribute 'incrementally' to a body of information that we all share. (Stokke and Fallis, 2017: 285) Stokke and Fallis point out that inquiry into large matters reaches its goal not by asking only one (big) question, but typically, by conducting 'subinquiries'. More often than not, it is these subinquiries that form the topics of a conversation – Stokke and Fallis call these 'questions under discussion' or 'QUD'. (2017: 286)

For Stokke and Fallis, bullshitting is indifference not to what one says in isolation, but to whether what one asserts advances a discussion by answering a QUD or not. With this view of what QUD's are in hand, they draw Frankfurt's distinction between lying and bullshitting as follows. Lies are direct (but false!) answers to QUD's. (2017: 288) By contrast, bullshitting is indifference to the QUD. (Stokke and Fallis, 2017: 279)

Stokke and Fallis's definition of bullshit then becomes:

A is bullshitting relative to a QUD  $q$  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. (Stokke and Fallis, 2017: 295)

Stokke and Fallis show how a bullshitter may be concerned with the truth (like the exam taker or the politician is when filibustering). They may care about the strict truth of what they say, they just do not care whether what they say is an answer to the question under discussion or not. Stokke and Fallis's account works particularly well for the Michael Howard account above – in the example it is clear that Howard speaks the literal truth, but evades the QUD.

### *Bullshit as not caring about knowledge*

Kimbrough (2006) and Fallis (2009: note 6) both suggest that bullshitting may amount to asserting something without having any evidence for it or without caring about the evidence. The idea has a certain appeal, in that making a bullshit statement is a tactic often employed



by people who have not taken the care to investigate what they are speaking about properly – the bullshitter often says something ‘without being able to back it up’. It is highly doubtful, however, that we can define bullshit as ‘asserting  $p$  without having evidence for  $p$ ’. While it may be imprudent, we often assert things without evidence. Take things that we assert on blind faith, like ‘God loves us’ or ‘everything will be all right in the end’; or take things said in hope like ‘I’m going to win the lottery tonight – I just know it’. Such things are often said without any evidence (even against better evidence!), yet, when said seriously they do not amount to bullshitting.

Gjelsvik presents a more sophisticated account. He holds that the bullshitter is not unconcerned with truth (as Frankfurt suggests) or evidence (as Kimbrough and Fallis suggest), but is unconcerned with knowledge. (forthcoming: 2) Starting from observations about the work of Kristeva, Gjelsvik examines whether the producer of academic bullshit is concerned with truth or not and what this concern means. Gjelsvik holds that the bullshitter may be concerned with truth (as Cohen has pointed out), but may not be concerned with *knowledge* in a more comprehensive sense. For Gjelsvik, a concern with knowledge implies that one cares about and tries to abide by the precise evidential norms that govern assertions in a particular field. He holds that anyone who truly cares about knowledge tries to speak as precisely as is necessary in a particular context so as not to say something that is inadvertently imprecise or misleading. By contrast, the bullshitter is apt to wade into a field that they do not know much about and makes reckless statements without regard for the precise state of knowledge or the exact evidential norms in effect in a context.. As an example of bullshitting, Gjelsvik (forthcoming: 11) mentions the way that Julia Kristeva makes big claims about set theory in her work on politics without, apparently, knowing much about set theory. Kristeva, for instance held that in Marx’s thought the state is a fiction that can be compared to the set of all sets. (Gjelsvik, forthcoming: 5)<sup>10</sup> Gjelsvik holds that, subjectively, or to herself, Kristeva may care about truth and knowledge – after all, she appears quite serious in her writing and has been said to be a very ‘scholarly’ writer. But objectively, Kristeva does not satisfy the conditions for knowing much about set theory, because she does not hold herself to the most obvious standards of evidence necessary to make claims about set theory (such as, one assumes, being able to formalise such claims). Finally, then, Gjelsvik’s proposal is that the bullshitter is not unconcerned with truth, but is unconcerned with knowledge. (2017: 5) As he puts it, one can

[take] oneself to be concerned with knowledge but actually [fail] to be properly so concerned. The latter state is possible because it is compatible with being in possession of the general concept of knowledge and caring about it, while failing to realize what knowledge amounts to in a field with specific, historically given requirements upon knowledge. (forthcoming: 10)

### **Defending Frankfurt’s account**

It is beyond the scope of this chapter to consider and criticise all of the alternative accounts of bullshit on the table. However, in a handbook of assertion, it is natural to elaborate and defend the assertion-based account of bullshit *par excellence* – that is, Frankfurt’s. I will therefore confine myself to offering brief criticisms of each of the alternative accounts in Frankfurtian vein in an effort to show that the leading account of bullshit is still Frankfurt’s.

---

<sup>10</sup> Gjelsvik points out that Kristeva makes similar handwaving gestures about Gödel’s incompleteness theorems.

Following that, I will further explain bullshit in Frankfurt's terms as a perversion of assertion distinct from lying.

### *Bullshit is not nonsense or obscurantism*

I start by considering two problems with the 'product' or 'output'-centred account of bullshit.

A first problem with Cohen's account of bullshit as 'unclarifiable unclarity' is identified by Frankfurt: Cohen does not define what it is for an utterance to be unclear, in fact, Cohen holds that one cannot define clarity (2002: 332). This, Frankfurt holds, amounts to Cohen '...hoisting [his] account of bullshit by its own petard...' (Frankfurt, 2002: 341 – 2) Second, it is doubtful whether we can indeed define bullshit simply as unclarity. Cohen holds that bullshit is simply 'nonsense'. (2006: 130) Nonsense, in the philosophical sense, is a sentence that, while it may appear meaningful, is in fact not meaningful. Here is an example: 'colourless green ideas sleep furiously'. (Chomsky, 1957) While this sentence pretends to say something about how the world is, it in fact says nothing – it has no content. If this is what nonsense is, the problem for Cohen is that no-one can truly believe it either: as no-one can have any idea of what it would be for colourless green ideas to sleep furiously, no-one can genuinely believe that colourless green ideas sleep furiously. But in order to explain what Cohen seeks to explain (essentially *honest* academics telling us things like that the state is like the set of all sets), Cohen has to assume that these people actually *do* believe what they say. To say something *honestly*, one has to believe what one says; but if bullshit is pure content-less non-sense, it cannot be believed and, therefore, also cannot be said honestly. (Kotzee, 2007)

### *Bullshit is not misleading implication*

In his account, Webber makes a number of powerful points about bullshit – in particular that it is possible to bullshit *by* asserting a truth. The observation that bullshit often *works* on the level of what is implied, rather than what is said, is surely right. However, it is important to establish to what extent this is a *criticism* of Frankfurt's view. In his paper, Webber admits that the phenomenon that Frankfurt identifies – 'bullshit assertion', that is asserting p while being indifferent to the truth of p – exists; he merely holds that bullshit can also function on the level of implication and that a proper assertion can be a vehicle for bullshit implication. Webber leaves Frankfurt's account of bullshit assertion largely alone and, instead, provides an additional analysis of bullshit conversational implicature. So we must ask: Can the implicature account displace the assertion account in the end? It seems to me not.

Frankfurt is careful to distinguish, on the level of assertion, between the lie and the bullshit assertion. Lying is saying what you believe to be not-true; bullshit assertion (to borrow Webber's terminology) is saying what you have no real belief about, either way. The two analogues on the level of implication would seem to be:

Misleading implication: believing p, but implying not-p

Bullshit implication: implying p while you have no belief either way whether p or not-p

Webber does not make clear exactly which one of these two he calls 'bullshit conversational implicature'. In fact, the example he uses (the Michael Howard example), looks more like a

case of misleading implication than bullshit implication. After all, it seems that Howard believed that it is not true that he did not threaten Lewis (otherwise, he surely would have felt himself entitled to say: 'I did not threaten Lewis'), but he clearly tried to imply that he did not threaten Lewis.

Acknowledging that the Howard case is a bit difficult to interpret, let us give a clear example of misleading implication and of bullshit implication.

Misleading implication: You email me at noon, asking whether I've posted your cheque. I reply 'My assistant posts all cheques after the office closes at 5pm.'

Say that my assistant does usually post all my cheques each day after the office closes at 5pm, but that, in this case, I do not intend to give him a cheque to post to you. However, in my reply, I clearly imply that I will be sending you your cheque at 5pm today. This is classic or ordinary misleading implicature. There is no bullshit involved, as evidenced by slippery language or a tricky attitude to the truth. I use straight talk to imply that I will send you your cheque, but do not actually believe that I will send it.

By contrast, take a case of bullshitting through implication.

Bullshit implication: In the Philosophy Department at Downstate University, a professor says to an undergraduate 'So, you've got a PhD offer from Upstate? At last year's APA, the rumours were flying about the place. You of all people should think very hard before accepting.'

Let us say that the professor did hear some rumours about Upstate at last year's APA, but knows full well that they are only rumours and has no opinion whether the rumours are true or false. In what the professor tells the student, however, the professor hints at some wild implications, like: (1) that it is unsafe for the student to go to Upstate (were the rumours of discrimination or harassment? and would this particular student be at risk?), (2) that Upstate's standing has taken a knock and (3) that it would be more prudent to go somewhere else. Whatever implication the student comes to believe or not is of no consequence to the professor. The student may come to believe any or all of implications (1) – (3) or none, as far as the professor is concerned. However, through their bullshit they have fulfilled a particular aim – talking Upstate down (and appearing a little superior by comparison). *Now that* would be bullshit implication.

In order to complete Frankfurt's analysis of bullshit assertion on the level of implication, we would have to say what distinguishes these two kinds of cases (misleading implication and bullshit implication) and, in my view, Frankfurt's distinction still works perfectly here: misleading implication would involve implying p when you believe that not-p and bullshit implication would involve implying p when you do not believe either p or not-p (and do not even particularly care whether p or not-p). This shows that, while it is a helpful advance on Frankfurt's account, Webber's implication account actually needs Frankfurt's analysis in terms of indifference to truth to be complete - it does not replace Frankfurt's original account.

*Bullshit is not evading the question under discussion*

Stokke and Fallis situate bullshit in the ebb and flow of conversation. They hold that bullshit is (in essence) not advancing the conversation by contributing to it constructively, but only to

pretend to contribute to it. Stokke and Fallis may be right that, given a strong enough grip on what the question under discussion is, we can identify acts of bullshitting by paying attention to whether someone attempts to answer the QUD seriously. But in contexts where it is not clear what the QUD is, we would not be able to identify misleading attempts to dodge the QUD, and, therefore, be unable to pin down acts of bullshitting. Take an ordinary, aimless conversation at the office first thing in the morning:

‘nice day, isn’t it?’  
‘nice but busy’  
‘oh, the end of term isn’t far off’  
‘have anything nice planned?’  
‘if I can finish this writing’.

In an exchange like this, it is hard to say what the question under discussion is and whether assertions like ‘busy’ or ‘if I can finish this writing’ are productive moves in the conversation or are mere bullshit.

A further difficulty is that one may bullshit even by asking a question, that is, one may bullshit exactly by manipulating the QUD to suit one’s bullshitting purpose. Take the asking of misleading rhetorical questions (a not unfamiliar tactic in academia). Lacan, for instance, asks:

can’t the formalization of mathematical logic, which is based only on writing, serve us in the analytic process in that what invisibly holds bodies is designated therein?

Lacan’s major scholarly contributions were to psychoanalysis and, in the piece of writing in question, Lacan was addressing the nature of memory. Whatever one makes of his work in psychoanalysis, surely directing his reader – through the asking of a question – to think about the formalization of mathematical logic as giving insight into the nature of memory is an example of bullshit misdirection. By *asking* about mathematical logic, one may say, Lacan attempted to bullshit by steering the question under discussion in an obscure direction. While Stokke and Fallis see bullshitting as evading the QUD, the QUD may *itself* be bullshit and sticking to the Stokke and Fallis view creates bizarre results when we are dealing with bullshit QUD’s. Imagine that Lacan’s misdirection succeeds and he, his colleagues and students start a conversation on something bogus like the link between mathematical logic and memory. Imagine that I join the conversation and say:

Well, if you really want to understand memory, you should forget about mathematical logic and start by drawing the standard analytic distinction between imagining and remembering.

My intervention would be a departure from or a refusal to engage with the QUD that Lacan himself set and would therefore risk being labelled bullshit on Stokke and Fallis’s account; however, in the example my intervention was meant exactly to cut off all this bullshit and drag the conversation back to sanity. The problem for Stokke and Fallis is that relativising bullshit to the questions posed in a particular conversation runs the risk of misclassifying the bullshit that takes place within conversations that are bullshit from start to finish.

*Bullshit is not indifference to knowledge*

We have already seen that it is too strong to call any assertion made without subjective evidence ‘bullshit’. We often say things in blind belief or hope, but can do so honestly (that is, without attempting either to mislead or to bullshit). In response, Gjelsvik proposes an account according to which it is not subjective evidence that makes the difference whether one bullshits or not, but objective knowledge. For Gjelsvik, bullshitting is saying something that one feels is subjectively justified, but, when measured against the external standard of what truly counts as knowledge in that domain, is not justified.

In response, notice firstly that it does seem possible to make a claim about an area and fail to live up to the external evidential standards in that area without bullshitting. I am not a physicist or an engineer, but I have heard of continuing efforts to generate electricity through cold fusion. I have also read about the Fleischmann/Pons experiment and about Low Energy Nuclear Reactions, but I am first to admit that what I know about cold fusion is little more than the product of some light internet surfing. Now that you know how little I truly know about cold fusion, imagine I seriously assert:

‘the SunCell is our best chance of cold fusion power generation in the next ten years’.

Am I necessarily bullshitting? Or might it be that I am simply poorly informed? What would help is knowing something about my intentions in talking up the SunCell. Was I trying to fool you? Then I was bullshitting. Or was I completely earnest? Then I was not bullshitting, but merely talking shit.<sup>11 12</sup>

Further, notice that one can *know* something and still bullshit. To see this, it is only necessary to extend the example of bullshitting by implication encountered above a little. Imagine that, in the example above, the Downstate professor is not any old academic but is Director of the Downstate Graduate Programme in Philosophy. Let us imagine that the professor knows absolutely everything that there is to know about the reputations of different graduate programmes in philosophy, about their relative placement records and about hiring trends in the profession. Imagine the professor tells the student:

Bullshit implication (extended): ‘So, you’ve got a PhD offer from Upstate? At last year’s APA, the rumours were flying about the place. What you should really think about is how this is going to affect their placement record. You know when the rumours broke about West Side University three years ago, their placement success went down 51%. You of all people should think very hard before accepting.’

In the example, it is not clear exactly what the professor implies about Upstate. For instance, is (s)he implying: That the rumours about Upstate are true or not? That the rumours about Upstate are the same as the rumours about West Side? That a PhD from Upstate is going to disadvantage a student at a future job interview because of the truth of these rumours, or because of the mere existence of the rumours? That there is something about this student, in particular, that makes Upstate a bad choice for them in the context of the rumours? That the student is (actually) so talented that they should be more ambitious than going to Upstate? That it is better to stay at Downstate? The point is that, in this example, the trouble with the professor’s statement is not a lack of knowledge about hiring and placement or a complete

---

<sup>11</sup> I draw my example from *Chemical and Engineering News* who lean to the view that SunCell is a clever hoax. <https://cen.acs.org/articles/94/i44/Cold-fusion-died-25-years.html>

<sup>12</sup> Arguably, Frankfurt gives an account of bullshitting, while Cohen gives an account of ‘shit’ (as in ‘shit-talk’). See above.

disregard for the norms of the field (philosophy). Rather than a lack of knowledge and a lack of fine attunement to discourse in the field, the professor has both in spades and uses this exactly *to bullshit* the poor student.

Gjelsvik might counter that the problem is not a lack of knowledge on the part of the professor, but a disregard for how philosophers should talk. Philosophers should be precise, informative and truthful in what they say, not vague and misleading (like the professor). One may say that while, abiding by the letter of the law regarding how one should speak in the philosophy department, the professor did not comply with the spirit. That is quite true, but the question is which one of the two proposals on the table – bullshit is cavalier reporting of belief (Frankfurt) and bullshit is a lack of genuine knowledge about a field (Gjelsvik) – best captures when someone does not abide by the ‘spirit’ of a form of discourse. The issue can be phrased like this. Is the problem with the Kristeva and Downstate cases that the bullshitter does not comport themselves in accordance with the evidential and discursive norms operative in their field? Or is the problem that they do not really care about what they say and therefore end up saying or implying something that is contentless, vague or confusing?

The second approach, Frankfurt’s, seems able to handle the Kristeva and Downstate cases very well. The heart of the bullshit in the Kristeva case is surely that, in her writings on the nature of the state, she said some things about sets that are imprecise, wrong and misleading in an effort not to advance understanding of sets, but in an effort to boost her views on politics. In the Downstate case, the problem is not how much the professor knows about or cares about graduate study and placement in philosophy, the problem is more basic – it is that they say something damaging to the student about Upstate, without knowing or caring whether any of it is true and dressing up what they say to appear like concern for the student.

### **Bullshit as fake assertion**

Above I reviewed and (briefly) considered four alternative accounts of bullshit. Each time, I held that Frankfurt’s original account – that bullshitting involves taking a cavalier attitude to the content of what one says, but pretending to be serious – comes closer to what we really mean by bullshit. In this section, I aim to elaborate and adjust Frankfurt’s account as the account that best contributes to the understanding of assertion. I hold that bullshit is a distinctive perversion of assertion next to lying and that considering lying and bullshit alongside ‘proper’ assertion helps us to sharpen up our conception of what assertion is.

I start with Gjelsvik account of bullshit as a failure of knowledge. While I hold that Gjelsvik’s account of bullshit as making assertions in contravention of external knowledge norms will not work, Gjelsvik does the field a great service by raising the issue of what norm of assertion it is that the bullshitter exactly flouts in bullshitting. This is very helpful to consider, because, as we have seen above, investigation of the norm of assertion is an important part of the field and there is substantial debate regarding what the norms of assertion exactly are: is it that the speaker should assert only what is true (Dummett, 1973; Wright, 1992; Weiner, 2005), only what they believe (Sayward, 1966; Hindriks, 2007; Bach, 2008), only what is justified (Lackey, 2007; Kvanvig, 2009), or only what they know (Williamson, 2000, DeRose, 2002; Hawthorne, 2005)? Bringing the question home to bullshit, we can ask: what norm of assertion, exactly is being flouted by the bullshitter?

We can swiftly dismiss the idea that bullshitting is contravening a truth norm for assertion. As Webber has explained (and as I have further outlined above), one can bullshit *by* asserting the truth. This can happen on the level of implication or (even) of straightforward assertion. Imagine that I am applying for a job in middle management with some big, faceless company. I do not know or care what the company makes or sells and merely want the job for the high salary. Before the interview, I read *Management for Dummies* and then pepper my conversation at the interview with truisms from it: ‘an organisation’s greatest asset is its people’, ‘failing to plan is planning to fail’, ‘you can’t improve what you can’t measure’, ‘brand perception is brand reality’, etc. All of these assertions are true (as far as they go!), but are said by me not because I believe them and want the hearer to believe them too, but because I want to fill the conversation with clever, but uninformative, cant. One could also say something false and still bullshit (e.g. ‘The Eskimos have a hundred words for snow’ or ‘Mussolini made the trains run on time’); but, equally, one could say something false and *not* bullshit (take any honest mistaken assertion). Abiding by and not abiding by the norm of truth does not line up well with the difference between properly asserting and bullshitting.<sup>13</sup>

Rather than a truth norm of assertion, I think that Frankfurt conceives bullshitting as breaking a belief norm of assertion. Consider again his characterisation of the bullshitter:

Her statement is grounded neither in a belief that it is true nor, as a lie must be, in a belief that it is not true. It is just this lack of connection to a concern with truth – this indifference to how things really are – that I regard as of the essence of bullshit. (Frankfurt 2005: 33 – 4)

In this passage, Frankfurt stresses the importance of a concern with the truth as the most important matter. However, note that he paints the bullshitter’s failing not as failing to say what is true, as such, but as not grounding what they assert in a serious belief about the truth. As Frankfurt sketches the matter here, the speaker who:

Asserts properly: believes that p is true and asserts p to bring the hearer to believe p

Lies: believes p and asserts not-p to bring the hearer to believe not-p

Bullshits: does not believe either p or not-p and asserts p regardless not to bring the hearer to believe p but to achieve some other goal (pertinently to impress the hearer in some way)

On Frankfurt’s view, the bullshitter contravenes a norm of belief for assertion; the bullshitter asserts something that they do not believe (in fact, they have no belief about it either way).

We have rejected the possibility that bullshitting is breaking a truth-norm for assertion out of hand and have seen that Frankfurt supports a belief-norm for assertion, but that still leaves two other possibilities on the table: that bullshitting is the breaking of a justification or a knowledge norm of assertion. Let us start with the justification norm. In the broader study of the norms of assertion, such a norm is advocated by, for instance, Lackey (2007) and

---

<sup>13</sup> Further support can be found in some of the literature on lying that suggests that lying, too, is not the breaking of a truth norm. One can obviously say something false without lying (take any honest mistake). One can also say something that is true and still lie – if one *thinks* that what one says is false, but, unbeknownst to one, it is true. (see Fallis, 2009).

Kvanvig (2009). Plausibly, Frankfurt's account of bullshit, lying and assertion can fit into a justification framework as follows. The speaker who:

Asserts properly: is justified to assert p and asserts p in order to provide evidence to the hearer that p

Lies: is justified to assert p, but asserts not-p in order to provide (bad!) evidence to the hearer that not-p

Bullshit: is not justified to assert either p or not-p, but asserts p in order to provide evidence to the hearer that p, regardless

It is obvious that the kind of justification in play here is *internal* justification, that is the justification as it appears to the speaker themselves. According to this kind of account, if a person knows full well that they are not justified by the evidence that is available to them to say something and, knowing this full well they *still* say it, they are bullshitting. Arguably, this picture of bullshit captures something important about the kind of bullshit that operates in charged political contexts, when, seemingly against all evidence, people are still apt to say things like: 'smoking doesn't cause cancer', 'HIV doesn't cause AIDS', 'there is no global warming', etc.

Relatedly, bullshitting may be contravening a knowledge norm for assertion. Amongst all of the possible norm-based accounts of assertion, the knowledge account arguably garners most support, from, for instance, Williamson (2000), DeRose (2004) and Hawthorne (2005). Stated in terms of a knowledge norm, the schema above would read as follows. The speaker who:

Asserts properly: knows p and asserts p in order to bring the hearer to know p

Lies: knows p and asserts not-p in order to bring the hearer not to know that p (alternatively: to believe something that is not knowledge, i.e. not-p)

Bullshits: does not know either p or not-p and asserts p regardless in order to make it the case that the hearer does not know p either

Which of these three - bullshitting as asserting without evidence, as asserting without knowledge and as asserting without believing - best reflects what is wrong with bullshit?

Let us start with bullshit as breaking the knowledge norm. What is so wrong about asserting something without *knowing* it? Williamson, for instance, writes as follows about asserting without knowledge:

asserting that P without knowing that P is doing something without having the authority to do it, like giving someone a command without having the authority to do so. (Williamson, 1996: 509)

He goes on:

such assertions are unacceptable because the speaker knows that he lacks the requisite knowledge, even though he has a reasonable belief. (Williamson, 1996: 511)



It is clear, especially from the second quotation, that the main problem Williamson sees with asserting *p* without knowing that *p* is asserting an otherwise reasonable belief with too little evidence. As an example, he cites cases like asserting that one's friend's lottery ticket did not win in advance of the result; Williamson's point is that even if the probability that the lottery ticket did not win is extremely high, it is still wrong to assert that the ticket did not win, as one does not quite *know* that the lottery ticket did not win. (Williamson, 1996: 502) In short, Williamson has in mind, under the category of 'assertions made without knowledge' assertions that are believed but not backed by sufficient evidence; he does not have in mind assertions that are not even believed (the kind of assertion Frankfurt calls 'bullshit').

There is considerable dispute regarding whether one must only say what one knows. We often say something like 'a change will do you good' or 'your back pain is probably nothing serious' without truly knowing it. Williamson holds that the apparent acceptability of such assertions in ordinary conversation shows only how weakly we tend to police the knowledge norm of assertion, not that such a norm is not in force. However, others (e.g. Weiner, 2005: 231) hold that the everyday acceptability of such innocent statements shows that it is not entirely improper to say something without full knowledge. Importantly, note that, whatever we say about the proper assertibility of careless statements like those above, no-one holds that saying something without knowing it is thereby automatically bullshitting: one can say something without knowing it due to being ignorant or naïve. The point is that asserting *p* without knowing *p* is a *broad* problem; it encompasses ignorant, wishful and naïve assertion as well as (possibly) bullshit assertion and advocates of the knowledge norm of assertion would need some way to distinguish between these different cases in order to give a more precise account of bullshitting as breaking a knowledge norm. Put simply, breaking a knowledge norm for assertion is too coarse-grained: it does not capture bullshit and only bullshit.

Exactly the same is the case for bullshitting as asserting something without evidence either way. Just as is the case with asserting without knowledge, asserting without evidence also encompasses a wide possible range of assertive failings, from ignorant, baseless and naïve assertion all the way through to bullshit assertion and advocates of an evidence norm would also need some kind of way to distinguish bullshit assertion from other kinds of un-evidenced assertion. Again, the evidence norm of assertion is too coarse to capture bullshit.

By contrast, asserting while not believing either way seems to capture bullshitting *and only* bullshitting quite well. Aside from cases like acting and joking (which are arguably not cases of assertion at all) the only thing to say of someone who says something without believing it either way is that they are bullshitting. This provides *prima facie* evidence that bullshitting is indeed the breaking of a belief norm.

Support for the idea that bullshitting is the breaking of a belief-norm of assertion (in particular: bullshit is asserting *p* without believing *p* or not-*p* either way) can be derived from a consideration of how or why we censure bullshitters. Catching a bullshitter about their business and censuring them for it is 'calling bullshit'. Let us pay attention, then, to calling bullshit: how does the one who calls bullshit formulate their charge against the bullshitter?

In my view, a clear example of bullshit was Bill Clinton's assertion during his Senate impeachment hearing that whether he had committed perjury during earlier legal proceedings depended on what the word 'is' means. At the impeachment hearing, Clinton was asked whether his lawyer Robert Bennett had made a false statement on his behalf during Paula

Jones's suit against Clinton for sexual harassment. Bennett had told the judge in the Paula Jones case that:

there is absolutely no sex of any kind in any manner, shape or form, with President Clinton [and Monica Lewinsky]

At his impeachment, Clinton was asked whether this was 'an utterly false statement' and he replied:

It depends on what the meaning of the word 'is' is.

Opinion divides as to Clinton's honesty in making this statement. Speaking narrowly, Clinton had a point. His lawyer was speaking at a time when sexual relations between Clinton and Ms Lewinsky had already ended, so, speaking in the present and not the past tense, what the lawyer said was strictly true. However, Clinton did have sexual relations with Lewinsky at an earlier point and attempted to hide this; as said in context, many would take the lawyer's vehement denial ('there is no sex') to cover not only the present state of affairs, but also the past. Moreover, it is not difficult to see a clear tactic at work in Clinton's assertion 'it depends what the meaning of the word 'is' is'. By answering like that, Clinton tried to dodge the real question – whether he and/or his lawyer lied – and tried to turn discussion to a narrow and technical debate about semantics, because such a digression may either mislead hearers into thinking he is innocent or at least disrupt or delay the flow of questions about his honesty.

As a point of historical fact, the Senate found Clinton not guilty of perjury, although later a Civil Court found him guilty of perjury. The score apparently reading 'Clinton 1 – Accusers 1', let us assume, for the purpose of the example, that we do not know for certain whether Clinton was asserting properly when he said 'it depends on what the meaning of the word 'is' is' or whether he was bullshitting.<sup>14</sup> The important question for our purposes is: how do we decide cases like these? When 'calling out' someone like Bill Clinton for bullshit are we wont to say:

You don't *know* that your earlier honesty depends on the meaning of the word 'is', Bill!

Or: You don't have any evidence that the meaning of the word 'is' settles the matter of your honesty, Bill!

Or would we say:

C'mon Bill, *you* don't even believe that. We're asking if you were honest and you say it depends on the meaning of the word 'is'. How can you expect us to believe it if you don't even believe it yourself?

No-one but Clinton knows what he was truly thinking at the time he made his infamous statement. However, imagine that we did. Imagine that we knew the state of Clinton's belief *or* evidence *or* knowledge from the God's-eye point of view to the effect that:

---

<sup>14</sup> Most people realise that Clinton was not *lying* when he said 'it depends on what the meaning of the word 'is' is'. He did not say something completely contrary to what he believes, so he was either asserting properly or bullshitting.

Belief: Clinton does not really believe that any ambiguity surrounding the meaning of 'is' absolves him from the charge of dishonesty *or*

Evidence: Clinton does not have good evidence that the ambiguity surrounding the meaning of 'is' absolves him from the charge of dishonesty *or*

Knowledge: Clinton does not know that the ambiguity in the meaning of the word 'is' absolves him from the charge of dishonesty.

Which one of these three options best captures the charge that he is bullshitting? If we knew from a God's eye point of view that Clinton did not really *believe* that the matter of his honesty turns on the meaning of the word 'is', then the charge of bullshit sticks very firmly. By contrast, if we know that he did not *know* it or that he did not have good *evidence* for it, then the charge does not stick as strongly because either of those two cases are compatible with the reading that Clinton still somehow *believed* that his honesty turns on the meaning of the word 'is'. The one who says something without knowledge or without evidence may be guilty of saying something that is ignorant, naïve or downright silly, but is not perforce bullshitting. However, the one who says something without even believing it themselves is clearly bullshitting. For this reason, I think the success of the charge of bullshit turns most naturally on what someone believes and that bullshitting is the breaking of a belief-norm or assertion.

The fact that one can treat lying and bullshitting as the breaking of a belief norm of assertion provides some (weak) evidence that the norm of the whole phenomenon of assertion is a belief-norm. After all, if it is possible to explain these two most obvious perversions of assertion as the breaking of a belief-norm, it is plausible that assertion proper can be explained as abiding by that norm. However, in this chapter, I only made an in-depth case that bullshitting is the breaking of a belief-norm; lying may yet be explained in other terms. Moreover, even if lying is also best explained as the breaking of a belief-norm, whether the norm of assertion *tout court* is a belief-norm will depend on how central the explanation of these two perversions of assertion are to understanding of the whole phenomenon of assertion.

## **Conclusion**

In this chapter, I explained and defended the classic account of bullshit assertion – Harry Frankfurt's original account that bullshitting consists in not caring about the truth of what one asserts. In his account, Frankfurt positioned bullshit as a distinct corruption or perversion of assertion. In so doing, Frankfurt not only defined bullshit for the first time, but also shed light on the nature of lying and on the nature of assertion proper. If bullshitting is not caring about the truth of what one says, then lying is saying what one believes is not true and asserting properly is saying what one does believe is true. In this chapter, I explained the main rival accounts of bullshit and, in particular, investigated what norm of assertion bullshitting breaks. *Pace* Frankfurt, I held that bullshitting is best explained as the breaking of a belief-norm for assertion.

## References

- Bach, K. 2008. 'Applying Pragmatics to Epistemology' *Philosophical Issues* 18, 68 – 88.
- Black, M. 1985. 'The Prevalence of Humbug' In: *The Prevalence of Humbug and Other Essays*. Ithaca, NY: Cornell University Press.
- Bok, S. 1978. *Lying: moral choice in private and public life*. Pantheon.
- Brandom, R. 1994. *Making it Explicit*. Cambridge, MA: Harvard University Press.
- Buekens, F. and Boudry. 2015. 'The Dark Side of the Loon: explaining the temptations of obscurantism' *Theoria* 81 (2), 126 – 42.
- Carson, T. 2010. *Lying and Deception*. Oxford: Oxford University Press.
- Chomsky, N. 1957. *Syntactic Structures*. De Gruyter.
- Cohen, G. 2006. 'Deeper into Bullshit' In: Hardcastle, G. and Reisch, G. (eds) *Bullshit and Philosophy: guaranteed to get perfect results every time*. Chicago: Open Court.
- Craig, E. 1990. *Knowledge and the State of Nature: an essay in conceptual synthesis*. Oxford: OUP.
- Davidson, D. 1984. 'Radical Interpretation' In: *Inquiries into Truth and Interpretation*
- DeRose, K. 2002. 'Assertion, Knowledge, and Context' *Philosophical Review* 111 (2), 167-204.
- Dudley, D. 2016. 'The Problem with Mussolini and His Trains'  
<https://www.citylab.com/transportation/2016/11/the-problem-with-mussolini-and-his-trains/507764/>
- Dummett, M. 1991. *The logical basis of metaphysics*, Cambridge, Mass.: Harvard University Press.
- Elster, J. 2011. 'Hard and Soft Obscurantism in the Humanities and Social Sciences' *Diogenes* 58 (1-2), 159 – 70.
- Fallis, D. 2009. 'What is Lying?' *The Journal of Philosophy* 106 (1), 29 – 56.
- Frankfurt, H. 2005. *On Bullshit*. Princeton, NJ: Princeton University Press.
- Grice, H.P. 1989. *Studies in the Ways of Words*. Cambridge, MA: Harvard University Press.
- Gjelsvik, O. Forthcoming 2018. 'Bullshit Production' In: Stokke, A. and Michaelson, E. *Lying: Language, Knowledge, Ethics, Politics*. Oxford: OUP.
- Hawthorne, J. 2005. *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Hindriks, F. 2007. 'The Status of the Knowledge Account of Assertion' *Linguistics and Philosophy* 30 (3), 393 – 406.

- Kimbrough, S. 2006. 'On Letting it Slide' In: Hardcastle, G. and Reisch, G. (eds) *Bullshit and Philosophy: guaranteed to get perfect results every time*. Chicago: Open Court.
- Kotzee, B. 2007. 'Our Vision and Our Mission: bullshit, assertion and belief', *South African Journal of Philosophy* 26 (2), 163 – 75.
- Kvanvig, J. 2009. 'Assertion, knowledge and lotteries', in P. Greenough & D. Pritchard (eds.), *Williamson on knowledge*, Oxford: Oxford University Press, pp. 140–60.
- Lackey, J. 2007. 'Norms of Assertion' *Nous* 41 (4), 594 – 626.
- Peirce, C. 1934. 'Judgment and assertion', in *Collected papers of Charles Sanders Peirce*, Boston: Harvard University Press, Vol. V, pp. 385–7.
- Pullum, G. 1991. *The Great Eskimo Vocabulary Hoax and Other Irreverent Essays on the Study of Language*. Chicago: University of Chicago Press.
- Sayward, C. 1966. 'Assertion and Belief' *Philosophical Studies* 17 (5), 74 – 8.
- Searle, J. 1969, *Speech acts: an essay in the philosophy of language*, Cambridge: Cambridge University Press.
- Stokke, A. and Fallis, D. 2017. 'Bullshitting, Lying and Indifference to Truth' *Ergo* 4 (10), 277 – 309.
- Webber, J. 2013. 'Liar!' *Analysis* 73 (4), 651 – 9.
- Weiner, M. 2005. 'Must we know what we say?' *The Philosophical Review*, 114 (2), 227–51.
- Williams, B. 2002. *Truth and Truthfulness*. Princeton, NJ: Princeton University Press.
- Williamson, T. 1996. 'Knowing and Asserting' *Philosophical Review* 105 (4), 489 – 523.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wreen, M. 2013. 'A P.S. on B.S.: some remarks on humbug and bullshit' *Metaphilosophy* 44 (1-2), 105 – 15.
- Wright, C. 1992. *Truth and Objectivity*. Cambridge, Mass: Harvard University Press.