Etica & Politica / Ethics & Politics, XIII, 2011, 2, pp. 201-214

Floridi on Disinformation

Don Fallis
University of Arizona
School of Information Resources & Library Science
fallis@email.arizona.edu

The government doesn't lie, it engages in disinformation.
George Carlin

KEYWORDS
Information, defection, misinformation, Floridi

1. *Introduction*

We cannot acquire, all by ourselves, all of the knowledge that we need to live our lives. We often have to rely on information that we receive from others (cf. Hume 1977 [1748], 74). For instance, we learn much of what we know about the world from books, newspapers, television, and the Internet.

However, when we get information from others, we have to consider the possibility that they might be trying to *deceive* us (cf. Hume 1977 [1748], 77). For instance, politicians, corporations, and governments (and sometimes even reporters themselves) have certainly used the media to deceive the public (cf. Jackson and Jamieson 2007). And in the very early days of the World Wide Web, Luciano Floridi (1996) already recognized the potential it had for misleading people (cf. Wachbroit 2000).

In order to avoid being misled by such information sources, it is helpful to know exactly what we are dealing with. In particular, what is *disinformation*? Floridi (1996, 2005, 2011) has dealt with this question several times, and his views have evolved over the years. In this essay, I describe the

three accounts of disinformation that Floridi has offered. But I argue that, while Floridi has been getting closer to an adequate account of disinformation, he has not yet been successful. Each of his accounts is too broad (i.e., it counts as disinformation things that clearly are not disinformation) and/or too narrow (i.e., it fails to count as disinformation things that clearly *are* disinformation).[1]

2. *False Information is Not Information*

Floridi is certainly the foremost *philosopher of information* in the world. In addition, while philosophers as far back as Plato have done work in this area, Floridi deserves credit for first identifying it as an important subdiscipline of philosophy. And one of his main tasks in this area has been to say what *information* is. Before we get to his accounts of disinformation, we first need to look briefly at Floridi's account of information. This is necessary because, for Floridi, disinformation is defined in terms of information.

Basically, information is "well-formed, meaningful and truthful data" (Floridi 2011, 80). "On weekends Luciano took the train in from Oxford to see Francesca" is an example of a piece of information.[2] By contrast, "Twas brillig, and the slithy toves did gyre and gimble in the wabe" and "Bill Clinton did not have sexual relations with that woman, Miss Lewinsky" do not count as information. (The first because it is not meaningful and the second because it is not truthful.)

The most controversial element of his account of information, and the element that is most relevant for our purposes here, is that information must be true. Several noted philosophers (e.g., Dretske 1981, 45-46, Grice 1989 [1987], 371, Frické 1997, 887-90) agree with Floridi that "false information" is a contradiction in terms. According to Fred Dretske (1983, 57), "false

[1] It may not be possible to give an account that captures everyone's intuitions about disinformation. In particular, there may be borderline cases that people disagree about. But I argue that Floridi's accounts do not even get some uncontroversial cases right.

[2] See Jonathan Harr's 2005 nonfiction book, *The Lost Painting*, for further information about Luciano and Francesca.

information, misinformation, and (grimace!) disinformation are not varieties of information—any more than a decoy duck is a kind of duck."

But there are other philosophers (e.g., Fox 1983, 157, Fetzer 2004a, Scarantino and Piccinini 2010) who argue that any meaningful data counts as information. Moreover, Scarantino and Piccinini (2010, 323-26) point out that computer scientists and cognitive scientists use the term *information* in a way that does not require that it be true. Similarly, when information scientists say that a library is full of information, they do not mean to be referring to just that subset of the collection that happens to be true.

My own preference is to say that information need not be true. When we get some stuff from an information source, we have definitely received some information; we just do not yet know whether or not it is true. But I am also inclined to think that the debate about the status of "false information" is (no pun intended) just a matter of semantics. So, for purposes of this essay, I will follow Floridi and talk about "semantic content" (i.e., meaningful data) when the stuff in question is not necessarily true.

3. *The Process of Information is Defective (1996)*

According to Floridi's (1996, 509) first account, "disinformation arises whenever the process of information is defective." Roughly speaking, disinformation occurs if semantic content is altered at some point in "its lifecycle (creation, storage, retrieval, updating)" in a way that makes it more likely that people will be misled. This account captures prototypical instances of disinformation, such as fraud, hoaxes, and government propaganda. But in addition, things like censorship are also disinformation on this account. In fact, Floridi (1996, 510) even claims that "each form of disinformation need not necessarily be intentional." Thus, honest mistakes, such as *The Chicago Tribune*'s erroneous report that "Dewey Defeats Truman" in 1948, are disinformation on this account.

While we certainly have to be concerned with any type of semantic content that is defective, this account of disinformation seems to be too broad. That is, there are examples of semantic content, that it counts as disinformation, but that are not disinformation. Most researchers (e.g.,

Fetzer 2004b, Jackson and Jamieson 2007) take the term *disinformation* to pick out, specifically, semantic content that is *intentionally misleading*.[3] In fact, the *American Heritage Dictionary* defines it as "deliberately misleading information" and the *Oxford English Dictionary* defines it as "the dissemination of deliberately false information." So, it seems that honest mistakes should not count as disinformation.

In any event, in order to avoid being misled by information sources, it is important for us to understand the various different types of defective semantic content. In particular, we should at least distinguish between semantic content that is accidentally defective (*misinformation*) and semantic content that is intended to be defective (*disinformation*). After all, the clues that suggest that someone is lying to us are different from the clues that suggest that she just does not know what she is talking about.

In addition, even though the manipulation of "the process of information" is intentional in the case of censorship, it should probably not count as disinformation either. Withholding information usually just keeps people in ignorance. For instance, several repressive regimes have recently censored (or even shut down) the Internet to keep their citizens ignorant of protests going on in their own or other countries. Admittedly, it is possible to actually deceive people by withholding information (e.g., so as to preserve false beliefs that would be overturned if people had access to the information). But even so, simply withholding information seems very different from actually creating or spreading disinformation.

4. *The Source is Aware of its Nature (2005)*

With his second account of disinformation, Floridi (2005, section 3.2.3) does attempt to characterize a more specific type of defective semantic content. He writes that "when *semantic content* is *false*, this is a case of *misinformation* (Fox [1983]). And if the source of misinformation is aware of its nature, one

[3] That is, the semantic content is likely to mislead the particular people that it is intended to mislead. Semantic content can still be disinformation even if there are all sorts of other people (e.g., the Amazing Randi's of the world) who are not likely to be misled.

may speak of *disinformation*, as when one says to the mechanic 'my husband forgot to turn the lights off.'"

Now, it is pretty clear that false semantic content can be disinformation even if the *immediate* source of the semantic content is not aware that it is false. For instance, prior to the Falklands War, the British Ministry of Defense convinced reporters that there were no plans for a Normandy-style invasion of the islands (cf. Jackson and Jamieson 2007, 99-100). Nevertheless, these reporters clearly passed along *disinformation* to the British public (and to the Argentineans) even though they themselves were not aware that it was false. So, I think that we have to assume that Floridi means that false semantic content is disinformation if the *original* source (in this case, the British Ministry of Defense) is aware that it is false.

However, even with this clarification, Floridi's 2005 account of disinformation is still too broad. Unless semantic content is likely to cause people to acquire a false belief, it is not all that epistemically dangerous, and it does not deserve to be called *disinformation*. As I argue below, Floridi's 2005 account does not even insure that disinformation is misleading, much less that it is intentionally misleading.

Even if a speaker is aware that what he is saying is false, the semantic content is not necessarily misleading. For instance, as Marc Antony was aware, the conspirators who assassinated *Julius Caesar* were not all honorable men. Even so, Antony was not spreading disinformation when he said *sarcastically* that the conspirators were "all honourable men." The Roman people were not likely to believe that the conspirators were honorable men on the basis of his statement. Similarly, people today are unlikely to believe the report in *The Onion* that "Al Gore Places Infant Son in Rocket to Escape Dying Planet."[4] Thus, this is not an example of disinformation either.

[4] Likewise, the plaque on Tanque Verde Road in Tucson, Arizona that describes the "Sand Trout (Salmo Harenatus). Endemic to the dry washes of Southern Arizona, this fish is able to withstand extreme heat and the absence of water. It has become adept at swimming around in the sand, feeding on scorpions, sand lions, and juvenile horned toads. In spring, breeding males become brightly colored, with red on the fins, belly, cheeks and lips. The fine spots on the sand trout's body help break up its outline, protecting it from predators, but making it vulnerable to passing automobiles" is not disinformation even though a few gullible pedestrians might believe it.

By saying something false, Antony was actually trying to convey to the Roman people something true (viz., that the conspirators were *not* honorable men). And *The Onion* was just making a joke. But even if a speaker actually intends to communicate something false, the semantic content is still not necessarily misleading. For example, in his 1875 novel, *The Way We Live Now*, Anthony Trollope writes that "when Sir Felix swore that a policeman was holding him while Crumb was beating him, no one believed him. In such cases the liar does not expect to be believed. He knows that his disgrace will be made public, and only hopes to be saved from the ignominy of declaring it with his own words." Unlike Antony and *The Onion*, Sir Felix was lying, but it is not clear that he was spreading disinformation as no one would be fooled by his statement.

False statements that are not intended to deceive are known as *bald-faced lies*. Such statements can certainly be dangerous. For instance, if Sir Felix had been a more respectable member of the gentry, people might have had to take his false accusation seriously, and an innocent policeman might have gotten into trouble. However, bald-faced lies are unlikely to cause anyone to acquire a false belief.

Admittedly, any false statement can potentially mislead someone. For instance, many first time readers of *The Onion* think (at least for a few moments) that it is a real newspaper. But honest mistakes can also mislead people even though they are not intended to. For instance, after reading the paper on the morning of November 3rd, many Chicago residents probably believed that Thomas Dewey had been elected President. Indeed, even (true) information can mislead people if they misread it or mishear it. So, the mere fact that there is some chance that someone will acquire a false belief from some semantic content is not enough to make it a piece of disinformation.

5. *It is Purposefully Conveyed to Mislead (2011)*

Most recently, Floridi (2011, 260) has argued that "misinformation is 'well-formed and meaningful data (i.e. semantic content) that is false.' 'Disinformation' is simply misinformation purposefully conveyed to mislead the receiver into believing that it is information." On this account,

disinformation is essentially the same thing as a *lie*. According to a standard philosophical account of lying, you lie if you say something that you believe to be false with the intent to deceive.[5]

This account of disinformation is very close to the dictionary definitions. In addition, James Fetzer (2004b, 231) has claimed that disinformation "should be viewed more or less on a par with acts of lying. Indeed, the parallel with lying appears to be fairly precise." In fact, this equivalence is also suggested by George Carlin's trenchant remark above and by the very title of Russ Kick's 2001 book, *You Are Being Lied to: The Disinformation Guide to Media Distortion, Historical Whitewashes and Cultural Myths*.

Unlike his previous account, Floridi's 2011 account rules out sarcastic remarks, jokes, and bald-faced lies. But it is probably still too broad. This account requires that the semantic content is intended to be misleading. But it does not insure that the semantic content *actually is* misleading.[6] For instance, imagine that eight year-old Billy has come to believe (from listening to his parents) that Paul Krugman and his ideas are a serious threat to our economy and our way of life. So, he decides to discredit Krugman by spreading some disinformation. With this in mind, he posts "Paul Krugman has cooties" on Krugman's blog at the *New York Times*. Billy is sophisticated enough to know that Krugman does not really have cooties (that there is no such thing), but he hopes to mislead Krugman's readers. Thus, on Floridi's 2011 account, Billy's post counts as disinformation. However, while Billy is certainly trying to create disinformation, he fails to do so because his post is not misleading at all. No one reading Krugman's blog is going to acquire the false belief that Krugman has cooties.

6. *Floridi's Most Recent Account is Too Narrow*

In addition to being too broad, Floridi's 2011 account of disinformation is also too narrow. That is, there are examples of semantic content, that it does

[5] On this account of lying, bald-faced lies are not really lies (just as decoy ducks are not really ducks).

[6] While lies have to be intended to mislead (on the standard philosophical account), they do not have to actually be misleading.

not count as disinformation, but that really are disinformation. In fact, as I argue below, it rules out four distinct types of semantic content that deserve to be called *disinformation*.

7. *Visual Disinformation*

Like Christopher Fox (1983, 75), Floridi (2011, 82) focuses on "declarative, semantic information." But images can also be used to convey information. For instance, a photograph or a map can show you that "the Bronx is up and the Battery's down."

Moreover, images can be used to convey semantic content that is false as well as semantic content that is true. For instance, during the 2004 Presidential campaign, a photograph appeared to show John Kerry and Jane Fonda sharing the stage at an anti-Vietnam war rally. But it was really a composite of two separate photographs taken at two separate events (cf. Farid 2009, 98). Also, in order to protect their intellectual property, many cartographers add a few features to their maps that do not really exist in the world (cf. Monmonier 1991, 49-51). If these non-existent features show up in another map of the same area, the cartographer has good evidence that her work has been copied.

While you have to use words in order to lie, these examples suggest that images by themselves can be disinformation. In addition, misleading images might even be more epistemically dangerous than misleading words. For instance, people generally take photographs to be more compelling evidence than mere testimony. When we trust testimony, we know that we are putting our faith in the person who produced it. By contrast, a photograph has evidential value independent of the intentions of the person who produced it (cf. Moran 2005, 8-11).

Now, Floridi might claim that he does not mean to suggest that visual information is not information or that visual disinformation is not disinformation. In fact, Floridi (2011, 84) does mention in passing that maps can count as information. But his focus on textual information certainly tends to underplay this important type of disinformation.

8. *True Disinformation*

Since disinformation is a subset of misinformation on Floridi's 2011 account, disinformation must be false. However, there may be some disinformation that is literally true. For instance, during the 2009 debate over healthcare legislation, several opponents claimed that the law mandated "death panels." The law did make provisions (a) for optional end-of-life counseling and (b) for a group that would do cost-benefit analyses of medical treatments (cf. Rutenberg and Calmes 2009). So, there is a sense in which the law did mandate "death panels." But of course, the law did not require forced euthanasia, which is the false conclusion many people drew and that the opponents of the legislation intended them to draw.[7] Similarly, a television commercial that pitted *Black Flag Roach Killer* against another leading brand misled viewers about the effectiveness of *Black Flag* without showing anything that was literally false. According to Thomas Carson (2002, 189), "the demonstration used roaches that had been bred to be resistant to the type of poison used by the competitor."

Even if one agrees with Floridi that information must be true, these examples suggest that there may be some disinformation that is information. In addition, true semantic content that is intended to be misleading might even be more epistemically dangerous than false semantic content (cf. Schauer and Zeckhauser 2009, 44-46). At the very least, such disinformation is likely to be more prevalent than disinformation that is actually false. There is less of a social stigma against misleading people than there is against out-and-out lying to them. Also, if you do not actually say anything false, it is more difficult to prove that you were trying to mislead anyone. So, would-be deceivers have some incentive to stick to the truth, if they can, while still deceiving.

9. *Side Effect Disinformation*

As I noted above, honest mistakes should not count as disinformation. But even so, there may be some disinformation that is not intended by the source

[7] Of course, some people did out-and-out lie about the law requiring forced euthanasia.

to be misleading (as Floridi's 2011 account requires). Before I give an example, however, it is helpful to note that someone can be aware that some state of affairs is likely to be a consequence of her action even though she does not intend to bring about that consequence. For instance, if a general orders the bombing of a military base and she is aware that some civilians living nearby will be killed, she still may not *intend* to kill civilians. If the death of the civilians is not a means to achieving her ends (e.g., taking out the military base), it is only a very unfortunate "side effect" of the bombing (cf. McIntyre 2009).

In a similar vein, someone might intentionally create semantic content that is misleading, but not intend that anyone actually be misled. The case of the cartographer who inserts small errors into her maps comes close to being an example of this kind of disinformation. She does not intend to deceive the vast majority of the people who use her maps. If legitimate users are misled, it is just an unfortunate side effect of her scheme to protect her intellectual property. However, the cartographer does at least intend to deceive potential intellectual property thieves.[8]

But there are others examples where there is no intent to deceive at all. For instance, researchers have put false semantic content into *Wikipedia* to see how long it takes to get corrected. Also, educators have created websites with false semantic content (e.g., about the plight of the Pacific Northwest Tree Octopus) in order to teach people how to distinguish accurate from inaccurate semantic content on the Internet (cf. Wachbroit 2000, 10). But neither of these groups intend to deceive anyone. It is not a means to their ends. In the case of the researchers, they just want to see whether or not people are deceived. And in the case of the educators, they probably hope that no one (especially their students) is deceived.

Now, the creation of false semantic content for purposes of research or education may be morally justified. But that does not mean that the semantic content is not disinformation. In fact, the dissemination of disinformation can be morally justified even if it *is* intended to deceive. For instance, it was presumably acceptable for the Allies to try to fool the

[8] Of course, she does not know for sure that there is anybody looking to steal her work, and actually hopes that there is no such person.

Germans and the Italians with false messages and documents during the Second World War (cf. Rankin 2008).

10. *Evolutionary Disinformation*

Finally, there may even be some disinformation that is not even foreseen by the source to be misleading. Before I give an example, however, it is helpful to note that someone (or something) can deceive even though she has no intent to deceive (cf. Skyrms 1990, 72-82). For instance, there are species of insects that have evolved to look like sticks or to look like leaves. These insects certainly do not form any intentions with regard to the beliefs of potential predators. For example, they do not *intend* that potential predators believe that they are flora rather than fauna because of what they look like. However, it is no *accident* that potential predators are fooled. The deceptive camouflage gives these insects an adaptive advantage. In particular, they are less likely to be eaten. Thus, the deceptive camouflage is more likely to continue into the future (because the genes that create it will be passed on to future generations).

Now, the appearance of an insect is not semantic content. Thus, such deceptive camouflage is not disinformation. But the same sort of thing can occur with misleading words (or images) as well as with misleading appearances. For instance, the story of the virgin birth is arguably an example of disinformation of this kind. Virgin births are impossible, or at least highly improbable (cf. Jurgensen and Southworth 2010).[9] But the people that tell this story believe that it is true. So, they do not intend to deceive anyone. However, as in the stick insect case, the spreading of this falsehood is reinforced by several actual benefits that come from the falsehood being believed. One such benefit was pointed out by David Hume (1977 [1748], 78). People tend to experience an "agreeable emotion," a sense of "surprise and wonder," when they hear that a miracle has occurred. As a result, the people describing the miracle to them can take "delight in exciting the admiration of others." Thus, the story of the virgin birth is more likely to

[9] If you believe that the story of the virgin birth is true, insert your favorite myth from another culture here.

be told again in the future. This kind of disinformation arguably has the purpose (or *telos*) of misleading people even though people are not misled on purpose (as Floridi's 2011 account requires).[10]

11. *Conclusion*

None of Floridi's three accounts quite capture the critical characteristic of disinformation: viz., that it is intentionally misleading. His most recent account comes closest to doing so. Unfortunately, this account of disinformation is too restrictive. In particular, it rules out visual disinformation, true disinformation, side effect disinformation, and evolutionary disinformation.[11] So, while Floridi has taken us a long way in the right direction, it is not yet clear that we have an adequate account of what disinformation is. Of course, philosophers have been trying to say what *knowledge* is for over two thousand years, and still have not gotten that right (cf. Feldman 2003). So, since the term *disinformation* was only coined about fifty years ago, and Floridi has only been trying to define it for fifteen, it is early days yet.

---

[10] Even in this case, there *could be* someone who intends people to be misled, such as an extremely powerful being who makes sure that those people who tell the story of the virgin birth are surreptitiously rewarded and encouraged. But as long as this being does not tell the story himself and did not come up with the story in the first place, the *source* of the semantic content does not intend it to be misleading. Similarly, we can imagine a powerful newspaper publisher who would like the public to believe falsely in the existence of a dangerous international conspiracy. But instead of instructing her editors and reporters to publish false semantic content about this conspiracy, she simply hires the relevant conspiracy theorists as editors and reporters and lets them publish whatever they like. Since the publisher does not control exactly what semantic content is published in the newspaper, it is the editors and reporters who are the source, and they believe that it is true and not at all misleading. Even so, it seems like her newspaper is publishing disinformation.

[11] Floridi's original 1996 account captures all four of the counter-examples to his most recent account. But as I have argued, that account is clearly too broad.

References

Carson, Thomas L. 2002. "Ethical Issues in Selling and Advertising." Pp. 186-205 in *The Blackwell Guide to Business Ethics*, ed. Norman E. Bowie. Malden, Massachusetts: Blackwell.
Dretske, Fred I. 1981. *Knowledge and the Flow of Information*. Cambridge: MIT.
Dretske, Fred I. 1983. "Précis of *Knowledge and the Flow of Information*." *Behavioral and Brain Sciences* 6:55-90.
Farid, Hany. 2009. "Digital Doctoring: Can We Trust Photographs?" Pp. 95-108 in *Deception*, ed. Brooke Harrington. Stanford: Stanford.
Feldman, Richard. 2003. *Epistemology*. Upper Saddle River, New Jersey: Prentice Hall.
Fetzer, James H. 2004a. "Information: Does It Have to Be True?" *Minds and Machines* 14:223-29.
Fetzer, James H. 2004b. "Disinformation: The Use of False Information." *Minds and Machines* 14:231-40.
Floridi, Luciano. 1996. "Brave.Net.World: The Internet as a Disinformation Superhighway?" *Electronic Library* 14:509-14. http://www.philosophyofinformation.net/publications/pdf/bnw.pdf
Floridi, Luciano. 2005. "Semantic Conceptions of Information." *Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/information-semantic/
Floridi, Luciano. 2011. *The Philosophy of Information*. Oxford: Oxford.
Fox, Christopher J. 1983. *Information and Misinformation*. Westport, Connecticut: Greenwood Press.
Frické, Martin. 1997. "Information Using Likeness Measures." *Journal of the American Society for Information Science* 48:882-92.
Grice, Paul. 1989. *Studies in the Way of Words*. Cambridge: Harvard.
Hume, David. 1977 [1748]. *An Enquiry Concerning Human Understanding*. Indianapolis: Hackett.
Jackson, Brooks and Kathleen H. Jamieson. 2007. *Unspun: Finding Facts in a World of Disinformation*. New York: Random House.

Jurgensen, Zachary and Jason Southworth. 2010. " Jesus, Mary and Hume: On the Possibility of the Virgin Birth." Pp. 11-23 in *Christmas - Philosophy for Everyone*, ed. Scott C. Lowe. Hoboken, New Jersey: Wiley-Blackwell.

McIntyre, Alison. 2009. "Doctrine of Double Effect." *Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/double-effect/

Monmonier, Mark. 1991. *How to Lie With Maps*. Chicago: University of Chicago.

Moran, Richard. 2005. "Getting Told and Being Believed." *Philosophers' Imprint* 5(5). http://www.philosophersimprint.org/005005/

Rankin, Nicholas. 2008. *A Genius for Deception*. New York: Oxford.

Rutenberg, Jim, and Jackie Calmes. 2009. "False 'Death Panel' Rumor has some Familiar Roots." *New York Times*. http://www.nytimes.com/2009/08/14/health/policy/14panel.html

Scarantino, Andrea and Gualtiero Piccinini. 2010. "Information without Truth." *Metaphilosophy* 41:313-30.

Schauer, Frederick and Richard Zeckhauser. 2009. "Paltering." Pp. 38-54 in *Deception*, ed. Brooke Harrington. Stanford: Stanford.

Skyrms, Brian. 2010. *Signals*. New York: Oxford.

Wachbroit, Robert. 2000. "Reliance and Reliability: The Problem of Information on the Internet." *Report From the Institute for Philosophy & Public Policy* 20(4):9-15.