
What Is Disinformation?

DON FALLIS

ABSTRACT

Prototypical instances of disinformation include deceptive advertising (in business and in politics), government propaganda, doctored photographs, forged documents, fake maps, internet frauds, fake websites, and manipulated Wikipedia entries. Disinformation can cause significant harm if people are misled by it. In order to address this critical threat to information quality, we first need to understand exactly what disinformation is. This paper surveys the various analyses of this concept that have been proposed by information scientists and philosophers (most notably, Luciano Floridi). It argues that these analyses are either too broad (that is, that they include things that are not disinformation), or too narrow (they exclude things that are disinformation), or both. Indeed, several of these analyses exclude important forms of disinformation, such as true disinformation, visual disinformation, side-effect disinformation, and adaptive disinformation. After considering the shortcomings of these analyses, the paper argues that disinformation is misleading information that has the function of misleading. Finally, in addition to responding to Floridi's claim that such a precise analysis of disinformation is not necessary, it briefly discusses how this analysis can help us develop techniques for detecting disinformation and policies for deterring its spread.

INTRODUCTION

According to Luciano Floridi (2011), *philosophy of information* is primarily concerned with "how information should be adequately created, processed, managed, and used" (p. 15). But as he (1996) notes, we also need to study what happens when "the process of information is defective" (p. 509). Inaccurate and misleading information can be extremely dangerous.

When people are misled about important topics, such as investment opportunities, medical treatments, or political candidates, it can cause serious emotional, financial, and even physical harm. Libraries and other information services are at the forefront of efforts to ensure that people have access to *quality information* instead.

Whether it results from an honest mistake, negligence, unconscious bias, or (as in the case of *disinformation*) intentional deception, inaccurate information (or *misinformation*) can mislead people.¹ But in the same way that acts of terrorism tend to be more troubling than natural disasters, disinformation is a particularly problematic form of misinformation because it is no accident that people are misled. Unlike an honest mistake, disinformation comes from someone who is actively engaged in an attempt to mislead (Fetzer, 2004; Piper, 2002, pp. 8–9). Moreover, in addition to directly causing harm, disinformation can harm people indirectly by eroding trust and thereby inhibiting our ability to effectively share information with one another. Thus, identifying strategies for dealing with the problem of disinformation is particularly pressing.

Disinformation is nothing new. Prototypical instances include deceptive advertising (in business and in politics), government propaganda, doctored photographs, forged documents, and fake maps. A standard example is Operation Bodyguard, a World War II *disinformation campaign* intended to conceal the planned location of the D-Day invasion. Among other deceptions, the Allies sent out fake radio transmissions and created fraudulent military reports in a successful attempt to convince the Germans that a large force in East Anglia was ready to attack Calais rather than Normandy (Farquhar, 2005, p. 72). However, disinformation seems to have become much more prevalent in recent years.

New information technologies are making it easier for people to create and disseminate inaccurate and misleading information (Hancock, 2007). Hackers have posted false stories on the websites of news services like *Yahoo! News* and the *New York Times* (Fiore & Francois, 2002). Also, investors have been duped by websites that “impersonate” the websites of reputable sources of information, such as *Bloomberg News* (Fowler, Franklin, & Hyde, 2001). Software now allows people to convincingly manipulate visual images (Farid, 2009). In addition, anyone with internet access can easily (and anonymously) insert inaccurate and misleading information into Wikipedia, “the free online encyclopedia that anyone can edit.” A notable example is when the entry on the journalist John Seigenthaler was edited to say that he was “directly involved in the Kennedy assassinations” (Fallis, 2008, p. 1665).

Libraries and other information services can easily end up being unwitting (and sometimes witting?) conduits for the spread of disinformation. In addition to disinformation that patrons may access over the internet, many library collections include government propaganda and historical

fabrications (Nesta & Blanke, 1991; Sowards, 1988). As a result, a number of information scientists have begun to address the problem of disinformation (Calvert, 2001; Fallis, 2009; Hernon, 1995; Karlova & Fisher, 2013; Lynch, 2001; Piper, 2002; Rubin & Conroy, 2012; Skinner & Martin, 2000; Walsh, 2010; Whitty, Buchanan, Joinson, & Meredith, 2012). In fact, the American Library Association (2005) has even issued a “Resolution on Disinformation, Media Manipulation, and the Destruction of Public Information.”

In order to effectively address the serious threat to *information quality* that disinformation poses, we need to develop techniques for identifying disinformation and policies for deterring its spread. However, in order to develop such techniques and policies, we first need to improve our understanding of the nature and scope of disinformation. Most importantly, we need to be able to distinguish disinformation from other forms of misinformation.² After all, the clues that someone is lying to us are likely to be different from the clues that she/he just does not know what she/he is talking about.

In this paper, I survey the analyses of the concept of disinformation that have been offered by various information scientists and philosophers (most notably, Floridi). I provide counterexamples to show that the existing analyses either exclude important forms of disinformation and/or include innocuous forms of information that should not be counted as disinformation. I then propose and defend a new analysis of the concept of disinformation. I close by briefly indicating how this analysis can help us deal with this serious threat to information quality.

CONCEPTUAL ANALYSIS

The standard philosophical methodology for clarifying concepts is known as *conceptual analysis* (Margolis & Laurence, 2011; Melnyk, 2008). The goal of conceptual analysis is to identify a concise set of necessary and sufficient conditions that correctly determines whether or not something falls under the concept in question. For example, an important project in philosophy is the analysis of the concept of *knowledge* (Feldman, 2003). The traditional analysis is that knowledge is *justified true belief*; in other words, the suggestion is that a) being justified, b) being true, and c) being a belief are necessary and sufficient for something to be a piece of knowledge. Similarly, an important project in information science is the analysis of the concept of *information* itself. Many different analyses of information have been proposed (Floridi, 2005; Fox, 1983, pp. 39–74).

Philosophers test a proposed analysis of a concept against a set of data points. More specifically, they check to see whether what the analysis says about various cases matches the intuitions of competent speakers of the language about those cases.³ There are two ways that an analysis might turn out to be inadequate. First, an analysis can be *too broad*; for instance,

an analysis of knowledge might count as knowledge some things that clearly are not pieces of knowledge. Gettier (1963) famously identified some justified true beliefs that we would definitely not want to classify as knowledge (Feldman, 2003, pp. 25–28). Second, an analysis can be *too narrow*; for instance, an analysis of knowledge might not count as knowledge some things that clearly are instances of knowledge.

With respect to the concept of disinformation, the goal is to find a list of criteria such that all instances of disinformation satisfy these criteria, and only instances of disinformation satisfy these criteria. In this paper, I offer cases that indicate that the existing analyses of disinformation are either too broad, too narrow, or both. I offer a new analysis that successfully handles all of these cases (or “data points”).

Using the intuitions of competent speakers of the language to test a proposed analysis is a good way to find out whether an analysis captures the *common usage* of terms like *knowledge* and *disinformation*.⁴ However, it is important to note that this is just a means to an end. We are ultimately concerned with what knowledge and disinformation really are, and not just with what people happen to think they are. As the philosopher John Austin (1956) explained, “We are looking again not *merely* at words . . . but also at the realities we use the words to talk about: we are using a sharpened awareness of words to sharpen our perception of, though not as the final arbiter of, the phenomena” (p. 8; emphasis in original). Thus, in addition to capturing our linguistic intuitions about disinformation, it matters whether the resulting analysis can actually help us to detect it and deter its spread. The ultimate test of an analysis of disinformation will be such practical applications.

THREE FEATURES OF DISINFORMATION

Disinformation Is Information

Before discussing the various analyses of disinformation that have been proposed, I begin by identifying three important features of disinformation that should be fairly uncontroversial. The first of these features is that disinformation is a type of *information*. Now, as noted above, several different analyses of information have been proposed. Therefore, exactly what it means to say that disinformation is a type of information depends on which analysis of information we choose to adopt. However, for our purposes here, we will not need to settle on a specific analysis of information. Almost everything that I will say about disinformation holds, regardless of the specific analysis of information we choose to adopt.

In this paper, I will primarily be appealing to just one central (and fairly uncontroversial) feature of information: that information is something that *represents* some part of the world as being a certain way. In other words, it is something that has *semantic* (or *representational*) content (Florida,

2011, p. 80; Scarantino & Piccinini, 2010, p. 324). For instance, the text “The cat is on the mat” represents the cat as actually being on the mat; also, a photograph of the cat can represent it as being on the mat.

Although I do not want to defend a specific analysis of information, there are two aspects of certain analyses that I am inclined to resist. First, according to some analyses of information (for example, Bates, 2006; Buckland, 1991), *any* object that “one might learn from,” including “fossils, footprints, and screams of terror,” counts as information. The information scientist Michael Buckland (1991) has most famously defended this notion of “information-as-thing.” I do not dispute his main claim that pieces of information are objects in the world rather than mental processes in the head. Prototypical instances of information, such as texts and photographs, are clearly objects. Moreover, I agree that many other objects, such as fossils and footprints, can be quite “informative.” Although such objects do not *represent* the world as being a certain way, they can certainly allow someone to *infer* that the world is that way. For instance, a footprint in the mud might tell Sherlock Holmes who the murderer is. However, as Buckland himself admits, “to include objects and events, as well as data and documents, as species of information is to adopt a broader concept than is common” (p. 356).

If we take disinformation to be a type of information on Buckland’s analysis of information, our analysis of disinformation is likely to be too broad. Any kind of deceptive activity would count as disinformation. For instance, in addition to fake radio transmissions, the Allies built fake tanks and airplanes out of rubber and canvas in their attempt to convince the Germans that the D-Day invasion would take place at Calais (Farquhar, 2005, p. 73). These objects would count as disinformation.⁵ As a result, we would really have no need for a special term for *disinformation*. (We could just talk about *deception*.) In order to avoid counting all deceptive activities as disinformation, I take the term *information* to refer just to what Buckland (1991) calls “representations” (p. 359).

Second, according to some analyses of information (Dretske, 1983, p. 57; Floridi, 2011, p. 260), something can only count as information if it is *true*. As the philosopher Fred Dretske (1981) points out, “information is, after all, a valuable commodity. We spend billions on its collection, storage, and retrieval. People are tortured in attempts to extract it from them. Thousands of lives depend on whether the enemy has it. Given all the fuss, it would be surprising indeed if information had nothing to do with truth” (p. 46). Thus, according to him (1983), “false information, misinformation, and (grimace!) disinformation are not varieties of information—any more than a decoy duck is a kind of duck” (p. 57).⁶ However, while the term *information* is sometimes used in a sense that implies truth, it is also commonly used in a sense that does not (Fox, 1983, p. 157; Scarantino & Piccinini, 2010, pp. 323–326). For instance, 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.

In any event, if we take disinformation to be a type of information on Dretske's analysis of information, our analysis of disinformation will be too narrow. Prototypical instances of disinformation, such as the Wikipedia entry claiming that Seigenthaler was involved in the Kennedy assassinations, clearly involve representing the world in a way that it is *not*. Thus, in this paper, I take the term *information* to refer to representational content that is false, as well as to representational content that is true.

Disinformation Is Misleading Information

The second important feature of disinformation is that it is a type of *misleading* information; that is, it is information that is *likely* to create *false beliefs*. It is this feature that makes disinformation dangerous and worthy of our concern. As noted above, disinformation can potentially have harmful consequences by virtue of misleading people; if a piece of information is not likely to create false beliefs, then it is not really a serious threat to information quality.

It should be noted that while a piece of disinformation must have the propensity to mislead, it does not have to actually mislead someone on any given occasion. Just like lying, disinformation is not a "success term." You are still lying even if the person that you intend to mislead does not believe what you say (see Mahon, 2008). Similarly, you are still disseminating disinformation even if your target does not happen to buy it. Now, if a piece of disinformation does not actually mislead someone on a particular occasion, then it is not going to have harmful consequences on that occasion. However, because it has the propensity to mislead, disinformation always puts people *at risk* of suffering harm (epistemic and otherwise). Therefore, it is still something that we would like to be able to identify and deter the spread of.

Disinformation Is Nonaccidentally Misleading Information

The third important feature of disinformation is that it must be *no accident* that it is misleading. It is this feature that distinguishes disinformation from more innocuous forms of misleading information, such as honest mistakes and overly subtle satire. Focusing on nonaccidentally misleading information puts us in a better position to detect this specific type of misleading information; for instance, much like lie-detection techniques, we can look for the intention to mislead instead of just looking for errors.

These three features of disinformation will be important in my evaluation of the existing analyses of disinformation. For instance, suppose that an analysis counts a piece of information as disinformation even though it is not misleading at all *or* it is only accidentally misleading. That will be *prima facie* evidence that the analysis is too broad. Conversely, suppose that an analysis *fails* to count a piece of information as disinformation

even though it is misleading, *and* it is no accident that it is misleading. That will be *prima facie* evidence that the analysis is too narrow.

PREVIOUS ANALYSES OF DISINFORMATION

Floridi (1996)

In one of the earliest discussions of the concept of *disinformation*, the philosopher Luciano Floridi (1996) claimed that “disinformation arises whenever the process of information is defective” (p. 509). However, this analysis is too broad. When someone makes an honest mistake like the *Chicago Tribune* reporting that “Dewey Defeats Truman,” something in the process is defective. But such *accidental falsehoods* clearly are not disinformation.

An honest mistake can certainly be misleading. For instance, after reading the paper on the morning of November 3rd, many Chicago residents probably believed that Thomas Dewey had been elected president. However, it is only an accident that they were misled. In such cases, the source of the information does not intend to mislead anyone (and does not benefit from people being misled). Also, as suggested above, indicators of accidentally inaccurate information are likely to be different from indicators of intentionally inaccurate information. Thus, an honest mistake should not be counted as disinformation.

Floridi (2005)

Several years later, Floridi (2005) claimed 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 may speak of *disinformation*” (§3.2.3; emphasis in original). In other words, disinformation is inaccurate information that the source knows to be inaccurate. This analysis repairs the shortcoming with Floridi’s 1996 analysis. It does not count accidental falsehoods as disinformation; when they ran the “Dewey Defeats Truman” story, the editors of the *Chicago Tribune* were not aware that the story was false.

However, Floridi’s 2005 analysis is also too broad. For instance, when you tell someone a joke or speak sarcastically, you are aware that what you are saying is false. But you are not spreading disinformation.⁷ Even though jokes and sarcastic comments are false, they are not misleading, nor are they intended to be misleading. The speaker expects her audience to also be aware that what she is saying is false. In addition, roadmaps are always inaccurate to some degree, and the cartographers are well-aware of this fact; for instance, if roads were really drawn to scale, they would be too small to see (Monmonier, 1991, p. 30). But, clearly, not all roadmaps are disinformation.

A story in the *Onion* claiming that “Al Gore Places Infant Son in Rocket to Escape Dying Planet” was also not disinformation, despite the fact that

the editors were aware that the former vice president did no such thing. Although a few credulous people who do not know that the *Onion* is satire might have been misled (at least briefly), it would have been totally accidental. The editors of the *Onion* did not intend to mislead anyone. Moreover, such inaccurate information did not pose the same risk of harm to the recipient that prototypical instances of disinformation do.

Fetzer (2004)

The philosopher James Fetzer (2004) claims 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” (p. 231). In other words, disinformation is a statement that the speaker believes to be false and that is intended to mislead. This analysis repairs the shortcoming with Floridi’s 2005 analysis; for instance, it does not count jokes and sarcastic comments as disinformation. Jokes and sarcastic comments are not lies because they are not intended to mislead (Mahon, 2008).

However, Fetzer’s analysis is also too broad. Someone who intends to spread disinformation with a lie might not succeed in doing so. For instance, suppose that the police ask about your friend’s whereabouts and that you want to mislead them about where he is. You believe that he is staying with his cousins outside the city, so you say to the police, “He is hidden in the cemetery.” However, without your knowledge, your friend has actually hidden himself in the cemetery. In this case, you have lied to the police (Mahon, 2008). (This is the opposite of an honest mistake. Instead of accidentally saying something false when you intend to say something true, you accidentally say something true when you intend to say something false.) But even though such *accidental truths* are lies, they are not disinformation because they are not actually misleading.

In addition, and more importantly, this analysis is too narrow. Lies are linguistic expressions, such as “a wolf is chasing my sheep!” (Mahon, 2008, §1.1). Thus, Fetzer’s analysis is limited to linguistic disinformation; it does not count deceptive *images* as disinformation. However, doctored photographs and falsified maps are clearly prototypical instances of disinformation. For instance, during the 2004 Presidential campaign, a photograph was circulated that appeared to show John Kerry and Jane Fonda sharing the stage at an anti-Vietnam War rally, while it was really a composite of two separate photographs taken at two separate events (Farid, 2009, p. 98). Also, during the cold war, the Soviets deliberately falsified maps in an attempt to fool their enemies about where important sites were located (Monmonier, 1991, pp. 115–118). It is no accident when people are misled by such *visual disinformation* because that is precisely what the source of the information intended.

Floridi (2011)

In his most recent discussion of the concept of *disinformation*, Floridi (2011) claims 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” (p. 260). In other words, disinformation is inaccurate information that the source intends to mislead the recipient. This is very close to several dictionary definitions of disinformation; for instance, the *Oxford English Dictionary* defines it as “the dissemination of deliberately false information.”

This analysis is along the same lines as Fetzer’s analysis, but it repairs the shortcomings with the latter’s: first, since Floridi explicitly requires that disinformation be false, accidental truths do not count as disinformation; second, although he tends to focus on *propositional* information in his work, Floridi (p. 84) allows that images and maps count as information. Thus, visual disinformation counts as disinformation in Floridi’s 2011 analysis.

However, his analysis is also too broad. Even if she/he says something that actually is inaccurate, someone who intends to spread disinformation still might not succeed in doing so. For instance, suppose a young child (from a conservative family) attempts to discredit Paul Krugman by posting “Paul Krugman has cooties!” on Krugman’s blog at the *New York Times*. This inaccurate statement about Krugman is not going to be believed by any readers of the blog; even though such *implausible lies* are (unrealistically) intended to be misleading, they are not disinformation because they are not actually misleading.⁸

In addition, and more importantly, this analysis is too narrow. Although prototypical instances of disinformation are inaccurate, disinformation can sometimes be *accurate*. For instance, when pursuers, who did not recognize him, asked Saint Athanasius “Where is the traitor Athanasius?” he replied, “Not far away.” As it was intended to do, Athanasius’s truthful statement misled his pursuers about his identity. 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 Carson (2010), “the demonstration used roaches that had been bred to be resistant to the type of poison used by the competitor” (p. 187). In addition, politicians often use *spin* to mislead the public—that is, they selectively emphasize only certain facts (pp. 57–58). Like prototypical instances of disinformation, such *true disinformation* is intentionally misleading and poses a similar risk of harm to the recipient.⁹

In fact, there is another respect in which Floridi’s 2011 analysis is too narrow. Although disinformation is always misleading, it is not always intended to mislead. For instance, inaccurate information has been intentionally placed on the internet for purposes of education and re-

search. A website advertising a town in Minnesota as a tropical paradise was created to teach people how to identify inaccurate information on the internet (Piper, 2002, p. 19). In addition, inaccurate information has been inserted into Wikipedia to see how long it takes to get corrected (Halavais, 2004).¹⁰ In these cases, while these educators and researchers certainly have to foresee that people might be misled by their inaccurate information, they do not intend that anybody actually be misled (nor do they benefit from people being misled). Thus, Floridi's 2011 analysis does not count this sort of inaccurate information as disinformation.

However, such *side-effect disinformation* probably should count as disinformation. Just as with prototypical instances of disinformation, it is no accident when people *are* misled. Even though the educators and researchers do not intend to mislead anyone, they do intend their inaccurate information to be misleading. A fake website would not be a very effective tool for teaching people how to identify inaccurate information on the internet if it was clear to everyone that it was a fake.

Fallis (2009)

According to the information scientist Don Fallis (2009), *disinformation* is "misleading information that is intended to be (or at least foreseen to be) misleading" (§5). This analysis repairs the shortcomings with Floridi's 2011 analysis. First, since Fallis explicitly requires that disinformation be misleading, implausible lies do not count as disinformation. Second, since Fallis does not require that disinformation be inaccurate, true disinformation counts as disinformation. Third, Fallis does not require that disinformation be intended to mislead; the source of the information merely has to foresee that it is likely to mislead. Although the educators and researchers described above do not intend to mislead anyone, they do foresee that some people may be misled. Thus, *side-effect disinformation* counts as disinformation in Fallis's analysis.

However, his analysis is too broad. In addition to *side-effect disinformation*, it also counts some subtle forms of humor as disinformation. For instance, a significant number of people are actually misled by the satirical stories published in the *Onion* (Fallon, 2012). In fact, stories from the *Onion* are sometimes picked up by serious news sources. For instance, Iran's state-run news agency reported the results of a (fake) Gallup poll that found that "77 percent of rural Caucasian voters . . . would much rather go to a baseball game or have a drink with Ahmadinejad than spend time with Obama." Moreover, since the editors of the *Onion* are clearly aware that this sort of thing goes on, they do foresee (even if they do not intend it) that at least some people will be misled by the stories they publish. But it is just an accident that people are misled.

Fallis's analysis can easily be modified, however, so that it does not count satire as disinformation. We can simply leave off the "foreseen to

be misleading” clause and say that disinformation is misleading information that is intended to be misleading. This modified analysis still counts side-effect disinformation as disinformation. For instance, as noted above, although educators do not intend to mislead anyone with their fake websites, they do intend these websites to be misleading.

There are also many cases that might seem to show that Fallis’s analysis is too narrow. Someone can clearly spread disinformation even if she/he does not intend, and does not foresee, that people will be misled. For instance, a press secretary might innocently pass along disinformation on behalf of her boss. But in that sort of case, there is *someone* (namely, the boss) who does intend that people be misled. Therefore, this sort of case does count as disinformation in Fallis’s analysis.

However, there are other cases that do indicate that Fallis’s analysis is too narrow. Even when a source of information does not intend to mislead anyone and does not foresee that anyone will be misled, it may be no accident that the information is misleading. For instance, many different species of animals give *fake* alarm calls (Skyrms, 2010, pp. 73–75). When animals give alarm calls, there is usually a predator in the vicinity. However, about 10 to 15 percent of the time, animals give alarm calls even when there is no imminent threat; in such cases, the call tends to cause other animals of the same species to run away and leave food behind, which the caller can then eat.

There is some evidence that primates understand that other primates can have false beliefs (Fitzpatrick, 2009). Thus, it could be that primates do intend to mislead conspecifics with their fake alarm calls. However, when less sophisticated animals who lack this understanding, such as birds and squirrels, give fake alarm calls, they do not intend to mislead anyone, and do not foresee that anyone will be misled, about a predator being nearby. Even so, these animals (or at least their genes) systematically benefit from giving such *deceptive signals*.¹¹ In other words, there is a mechanism that reinforces the dissemination of this sort of misleading information—namely, the caller gets a free meal. Thus, like prototypical instances of disinformation, such *adaptive disinformation* is not misleading by accident and poses a similar risk of harm to the recipient.

If Fallis’s analysis simply failed to count the deceptive signals of *animals* as disinformation, it might not be a serious objection; after all, it is not clear that alarm calls have representational content and thus count as information. However, humans can also disseminate adaptive disinformation. For instance, many of the people who disseminate conspiracy theories (for example, that the president was not born in the United States or that the US government was behind the 9/11 terrorist attacks) believe that what they are saying is true. Thus, they do not intend to mislead anyone, or foresee that anyone will be misled, by what they say. Even so, just as with prototypical instances of disinformation, these false claims can

mislead people, *and* it is no accident that people are misled. There is a mechanism that reinforces the dissemination of these false claims. By promoting such false claims, certain websites and media outlets are able to attract more readers and viewers who find these claims convincing.¹²

Skyrms (2010)

Recent work in biology on deceptive signaling in animals might provide an analysis of the concept of *disinformation*. According to the philosopher Brian Skyrms (2010), "If misinformation is sent systematically and benefits the sender at the expense of the receiver, we will not shrink from following the biological literature in calling it *deception*" (p. 80). Although Skyrms and the biologists he cites use the term *deceptive signal* rather than *disinformation*, they are trying to capture essentially the same concept. Thus, we might say that disinformation is misleading information that systematically benefits the source at the expense of the recipient. This analysis repairs the shortcoming with Fallis's analysis. Although animals that give fake alarm calls may not intend to mislead others, they do systematically benefit from others being misled. Thus, adaptive disinformation counts as disinformation on Skyrms's analysis.

However, his analysis is too narrow. Most of the time, disinformation imposes a cost on the recipient, as when the villagers waste their time running to the shepherd boy's aid. However, disinformation need not always impose a cost on the recipient; in fact, it is sometimes intended to benefit the recipient. For instance, when a friend asks you how he or she looks, you might very well say "You look great!" in order to spare his or her feelings, even if it is not true. Admittedly, such *altruistic disinformation* does not pose the same risk of harm to the recipient that prototypical instances of disinformation do; but like prototypical instances, altruistic disinformation can be intentionally misleading.

Skyrms's analysis can easily be modified, however, so that it counts altruistic disinformation as disinformation. We can simply leave off the "at the expense of the recipient" clause and say that disinformation is misleading information that systematically benefits the source. It is really just the "systematic benefit to the source" clause that is needed to count adaptive disinformation as disinformation. In fact, Skyrms (p. 76) himself notes that his analysis might be modified in this way. He just failed to see that this sort of modification was actually necessary.

However, Skyrms's analysis is still too narrow because it rules out the possibility of disinformation that does not benefit the source. Most of the time, disinformation does systematically benefit the source; it often keeps the source out of trouble, as when Athanasius was able to escape his pursuers. Or it simply makes a social situation less awkward, as when you say to your friend "You look great!" Either way, the deceptive behavior is reinforced. However, disinformation need not always benefit the source; for

instance, in order to avoid embarrassment, people often lie to their doctors about their diet, about how much they exercise, or about what medications they are taking (Reddy, 2013). If their doctors are misled, it can lead to incorrect treatment recommendations that can harm the patient. (Patients who lie may not be aware of the risk they are taking, or they may not be able to help themselves even though they are aware of the risk.) Admittedly, this particular example of *detrimental disinformation* may not pose the same risk of harm *to the recipient* that prototypical instances of disinformation do. But as with prototypical instances, it is no accident that people are misled by such disinformation.¹³

A NEW ANALYSIS OF DISINFORMATION

Even though the modified Fallis analysis (in terms of an intention to be misleading) and the modified Skyrms analysis (in terms of a systematic benefit from being misleading) are too narrow, together they arguably capture all instances of disinformation. It would be unfortunate, however, if we had to resort to such a *disjunctive* analysis of disinformation. If an analysis requires two independent criteria, it suggests that we are really dealing with two separate phenomena (Kingsbury & McKeown-Green, 2009, pp. 578–581). However, there is something that unifies all of the cases of disinformation discussed above: disinformation is misleading information that has the *function* of misleading someone.

Roughly speaking, a *function* is “the action for which a person or thing is particularly fitted or employed” (American Heritage, 2000).¹⁴ For instance, the function of a heart is to pump blood; also, the function of a chair is to be sat on. According to this analysis, the distinguishing feature of disinformation is that its function is to mislead people.

It should be noted that there are at least two different ways that something might acquire a function (Graham, 2010, pp. 153–155; Krohs & Kroes, 2009). For instance, a heart has the function of pumping blood because that is what it *evolved* to do; by contrast, a chair has the function of being sat on because that is what it was *designed* to do. (Similarly, an *artificial* heart has the function of pumping blood because that is what it was designed to do.) In other words, the designer of the artifact *intended* it to have that function.¹⁵

Disinformation can acquire the function of misleading people in either of these two ways. Most forms of disinformation, such as lies and propaganda, are misleading because the source intends the information to be misleading. But other forms of disinformation, such as conspiracy theories and fake alarm calls, are misleading simply because the source systematically benefits from their being misleading. Even though they might differ in terms of how that function was acquired, all instances of disinformation are unified by the fact that they have a certain function. And however that function was acquired, it is no accident that the information is misleading.

This analysis of disinformation seems to include all cases of disinformation. For instance, the adaptive disinformation that caused a problem for the modified Fallis analysis has the function of misleading because the source systematically benefits from it being misleading. Also, the detrimental disinformation that caused a problem for the modified Skyrms analysis has the function of misleading because the source intends it to be misleading.¹⁶

In addition, this analysis seems to exclude those cases that are clearly not disinformation; for instance, accidental truths and implausible lies are not misleading. Also, while accidental falsehoods and satire can sometimes be misleading, the source of the information does not intend to mislead people, nor does she/he systematically benefit from people being misled. If people are misled, it is simply an accident. (See tables 1 and 2.)

A few things should be noted about this analysis of disinformation. First, the appeal to *functionality* makes disinformation a “historical” concept; that is, whether or not a piece of information counts as disinformation does not just depend on its “perceptible and potentially perceptible features” (Rowe, 2012, p. 157). In particular, it is not sufficient *that* the information is misleading; it depends on *how* the information came to be misleading. In this regard, the property of being disinformation is like the property of being a genuine Van Gogh. Even if it were possible to create a molecule-for-molecule reproduction of *The Starry Night*, it would not be a Van Gogh because it was not actually painted by Vincent Van Gogh (p. 173).

Second, the fact that different types of disinformation have acquired the function of misleading people in different ways can make a difference on how we might detect disinformation; for instance, we might look for evidence that someone has the intention to mislead.¹⁷ However, we might also look for evidence that someone benefits from people being misled.

Third, whether or not a piece of information counts as disinformation on this analysis depends on who receives it. This is because different things are misleading to different people. For instance, a piece of information that is misleading to Dr. John Watson may not be misleading to Sherlock Holmes, Esq. A prime example is a letter forged by Professor Moriarty that calls Watson back to their hotel in Meiringen while Holmes continues on alone to that “fearful place” known as the “falls of Reichenbach.” Although Holmes lets Watson go, the letter—supposedly written by an English lady in “the last stage of consumption” who wants to see an English doctor—does not fool him for an instant. However, if we want the concept of disinformation to be person-independent, we could utilize a “reasonable person” standard, along the lines of the Federal Trade Commission’s definition of deceptive advertising (Carson, 2010, p. 187; Fallis, 2009). In other words, we might require that disinformation be misleading to (and have the function of misleading) a reasonable person.

Table 1. Data points

Clearly disinformation	Reason	Clearly not disinformation	Reason
malicious lies (ML)	nonaccidentally misleading	truthful statements (TS)	not misleading
visual disinformation (VD)	nonaccidentally misleading	accidental falsehoods (AF)	only accidentally misleading
true disinformation (TD)	nonaccidentally misleading	jokes (J)	not misleading
Side-effect disinformation (SE)	nonaccidentally misleading	sarcastic comments (SC)	not misleading
adaptive disinformation (AD)	nonaccidentally misleading	accidental truths (AT)	not misleading
altruistic disinformation (AL)	nonaccidentally misleading	implausible lies (IL)	not misleading
detrimental disinformation (DE)	nonaccidentally misleading	satire (S)	only accidentally misleading

Fourth, the fact that a piece of information has the function of misleading does not mean that it had that function as soon as it was created or first disseminated. Although the satirical stories in the *Onion* are not intended to mislead people, such a story would become disinformation if the editors seriously defended its accuracy in the mainstream press. Similarly, a stapler does not have the function of holding a door open, but it can acquire that function if I start using it as a doorstop. In fact, a piece of information might even become disinformation through an act of omission (Carson, 2010, p. 56). If the source of a piece of information sees that people are misled, she/he will often have an obligation to correct their mistake. Therefore, if she/he has an opportunity to set the record straight and fails to do so, it is no longer an accident that people are misled.

Fifth, saying that disinformation has *the* function of misleading people is not meant to suggest that misleading people is the *ultimate purpose* of the dissemination of the information. Misleading people can be (and usually will be) just an intermediate step toward a further end. For instance, Athanasius only misled his pursuers in order to escape; even so, the function of his misleading utterance was to mislead them.

Finally, it should be noted that while the ultimate purpose of disseminating disinformation is often problematic, it is sometimes a good thing. For instance, the Allies were presumably justified in using disinformation to prevent Nazi world domination. Also, the hacker who placed false stories on *Yahoo! News* apparently did so in order to expose the security

Table 2. Analyses of disinformation and data points

	Clearly disinformation							Clearly not disinformation						
	ML	VD	TD	SE	AD	AL	DE	TS	AF	J	SC	AT	IL	S
Floridi (1996)	√	√	√	√	√	√	√		√			√	√	
Floridi (2005)	√	√		√		√	√			√	√		√	√
Fetzer (2004)	√					√	√					√	√	
Floridi (2011)	√	√				√	√						√	
Fallis (2009)	√	√	√	√		√	√							√
Fallis (modified)	√	√	√	√		√	√							
Skyrms (2010)	√	√	√	√	√									
Skyrms (modified)	√	√	√	√	√	√	√							
New proposal	√	√	√	√	√	√	√							

weaknesses of the website so that they could be fixed. In fact, disinformation can sometimes even be a good thing, *epistemically speaking*. Even if a piece of disinformation leads to a false belief, it might lead to more true beliefs in the long run. For instance, as Pratchett, Stewart, and Cohen (1999) point out, “a *lie-to-children* is a statement which is false, but which nevertheless leads the child’s mind towards a more accurate explanation, one that the child will only be able to appreciate if it has been primed with the lie” (p. 38). An example is telling a student that electrons orbit the nucleus of an atom like a little solar system.

FLORIDI’S WORRIES ABOUT THE ANALYSIS

In response to the counterexamples offered in this paper, defenders of the existing analyses of disinformation might try to bite these bullets; that is, despite the compelling fact that a piece of information is not misleading or is only accidentally misleading, they might insist that it nevertheless counts as disinformation. Alternatively, despite the fact that a piece of information is nonaccidentally misleading, they might insist that it does not count as disinformation. However, these seem like fairly big bullets to bite.

Instead of biting these bullets, Floridi (2012, pp. 306–307) has taken a different tack. Fallis (2011) presented several of these counterexamples to Floridi’s own analyses of disinformation. In his response, Floridi pointed out that there may not be a concise set of necessary and sufficient conditions that perfectly captures the concept of disinformation. I am happy to concede this point. Disinformation may be a “prototype concept”; that is, there may simply be prototypical instances of disinformation, with different things falling closer to or further from these prototypes. But even

if it is impossible to find an analysis that perfectly captures the concept of disinformation (that is, one that is “counterexample proof”), we can still strive to get as close as we can. It seems likely that doing so would help us to develop better techniques for identifying disinformation and better policies for deterring its spread.

However, in his response, Floridi also claims that there is no point in getting closer to capturing the concept of disinformation than he has already done. As he puts it, “The important question to ask is not how far we can go with the fine slicing of a variety of different meanings and peculiar cases, but whether this is worth the effort.” Floridi asks us to consider the following analogy: Suppose that someone wants to know “how many steel rails her company needs to buy to cover the distance between two cities.” With this purpose in mind, measuring the steel to the closest centimeter would be “rather pointless.” Measuring to the closest meter would clearly be sufficiently precise. According to Floridi, “the criticisms offered by Fallis in terms of visual disinformation, true disinformation, side effect disinformation, and [adaptive] disinformation are interesting, but they concern centimetres, whereas I was giving the length of the steel rail in metres.”

It is certainly true that the method of conceptual analysis has sometimes been taken to unnecessary extremes. For instance, we might question whether the various analyses of knowledge that have been proposed in response to Gettier’s (1963) counterexamples are really getting us closer to something with theoretical or practical utility. In fact, Weatherston (2003) has argued that despite the counterexamples, the justified-true-belief analysis of knowledge is good enough for philosophical purposes. I do not think this to be the case here, however; there are reasons to believe that none of the existing analyses of disinformation is good enough for dealing with this serious threat to information quality. The counterexamples that I have offered to these analyses are not just logical possibilities (or “peculiar cases”); these forms of disinformation are fairly common real-world phenomena that can be quite dangerous.¹⁸

Misleading images (that is, visual disinformation) might easily 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 (Moran, 2005, pp. 8–11). Moreover, visual disinformation may very well be more prevalent than linguistic disinformation; in fact, according to Mark Twain (1900/2002), “By examination and mathematical computation I find that the proportion of the spoken lie to the other varieties is as 1 to 22,894” (p. 169).

In a similar vein, Schauer and Zeckhauser (2009, pp. 44–46) have ar-

gued that misleading information that is true (that is, *true* disinformation) can be more epistemically dangerous than misleading information that is false. Moreover, true disinformation may very well be more prevalent than disinformation that is actually false. There is less of a social stigma against simply 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. Therefore, would-be deceivers have an incentive to stick to the truth, if they can, while still deceiving.

Finally, even if disinformation that is intended to mislead happens to be more prevalent than side-effect disinformation (that is not intended to mislead) or adaptive disinformation (that is not intended to mislead), it is still important to be aware of these forms of disinformation. For one thing, disinformation that is intended to mislead frequently is side-effect disinformation and/or adaptive disinformation. For instance, even when disinformation is only intended to mislead one particular person, it typically misleads many other people as a side effect. As Mearsheimer (2011) points out, "Leaders engaged in inter-state lying usually end up deceiving their own people, although they are not the intended audience" (p. 21). Also, even when disinformation is intended to mislead, there is often some additional mechanism that reinforces its dissemination. Although there are certainly exceptions (such as Pinocchio), liars tend to benefit in some way from telling lies.

SOME BORDERLINE CASES

Visual disinformation, true disinformation, side-effect disinformation, and adaptive disinformation are clearly important forms of disinformation. But it should be noted that there are other cases where it is not quite so clear that a piece of information has the function of misleading people and thus should count as disinformation. Since such borderline cases of disinformation could make a difference in how we try to detect disinformation or deter its spread, they deserve some consideration.

Content Is Not Misleading?

What if a piece of information is misleading (that is, it is likely to create false beliefs), but is not misleading about the accuracy of its content? For instance, suppose that you create a map of South America so that it looks like it was drawn by Europeans in the seventeenth century (for example, by its being on old parchment, with ornate lettering) to try to misled people into thinking that Machu Picchu was discovered in the seventeenth century. This is certainly information that has the function of misleading. Thus, your fake map counts as disinformation in my new analysis. However, the *content* of this map is completely accurate; it is not going to mislead anyone about the location of Machu Picchu. With this sort of case in mind, we might want to amend my new analysis and reserve the term

disinformation for information whose *representational content* has the function of misleading.

Not Very Misleading at All?

What if a piece of information creates a false belief just as it was intended to, but it was not at all *likely* to create a false belief? For instance, suppose that the young child's blog post, against all odds, does mislead someone into believing that Krugman has cooties. (After all, as the case of the *Onion* shows, people can sometimes fall for some pretty implausible stuff.) Although it would still be a stretch to say that the blog post was misleading (much less, that it had the function of misleading), it would not have been an accident that someone was misled. Similarly, suppose that a basketball player throws up a last-second shot from mid-court and it happens to go in. Although she or he got really lucky, the basket was not an accident; the player *was aiming* to put the ball through the hoop. It was not as if she/he just threw the ball up into the air at random and it happened to go in.

In my new analysis of disinformation, I have not specified exactly *how likely* it must be that a piece of information will cause false beliefs in order for it to count as misleading. In a similar vein, epistemologists rarely specify the exact degree of justification that is required for a belief to count as knowledge. Most simply say that "very strong justification" is required (Feldman, 2003, p. 21). Along the same lines, we might require that a piece of information be *very likely* to create false beliefs in order for it to count as misleading. Alternatively, we might adopt a more expansive account that simply requires that a piece of information make the creation of false beliefs more likely than chance.¹⁹

Misleading in Some Other Way?

What if a piece of information is misleading, but not in the way it was intended to be misleading? For instance, suppose that you want to mislead your friend into believing that Bisbee is southeast of Tombstone by drawing her a map. Since (unbeknown to you) Bisbee really is southeast of Tombstone, your map is not misleading with respect to the relative locations of the two towns. However, because you drew all of the roads with the same thickness, she does acquire the false belief that Tombstone and Bisbee are connected by an interstate highway (when, in fact, they are connected by a small country road). Even though your map is misleading, it does not have the function of misleading someone (about the size of the connecting road or anything else). Therefore, it does not count as disinformation in my new analysis. This seems like the correct result, since it was just an accident that your friend was misled about the size of the connecting road.

However, it should be noted that a piece of information might legitimately count as disinformation even if it is not misleading in exactly the

way it was intended to be. For instance, suppose that Moriarty creates a document cleverly designed to mislead Holmes about his current whereabouts. But it turns out that the document is not clever enough to fool someone with Holmes's vaunted powers of ratiocination. While the document *is* misleading to Watson, that is not what Moriarty intended. (Even if Watson had seen through the ruse and figured out the truth about Moriarty's whereabouts, it would not have mattered to Moriarty. He confidently expects that Holmes will arrogantly ignore anything that Watson claims to have deduced.) Even so, it is no accident that Watson is misled. Since Moriarty intends the document to be misleading to Holmes, he does intend it to be misleading to members of the species to which Watson belongs.²⁰ Thus, this misleading information does have the function of misleading people such as Watson.

Just Keeps People in the Dark?

What if a piece of information is not misleading (that is, it is not likely to create false beliefs), but is likely to keep people in ignorance? For instance, suppose that the Germans were extremely skeptical of any communications that they intercepted (which may very well have been the case). In that case, the fake radio transmissions would not have been likely to convince the Germans that the D-Day invasion would take place at Calais. Even so, the Allies might still have sent the transmissions in order to keep the Germans in the dark about where the attack would come (that is, to keep them from acquiring a true belief that the invasion would take place at Normandy). But in that case, the transmissions would not have been, strictly speaking, misleading. Therefore, they would not have counted as disinformation in my new analysis.

However, the fake radio transmissions would still have left the Germans epistemically worse off than they might have been (see Chisholm & Feehan, 1977, p. 144). This sort of manipulation of another person's beliefs can be just as dangerous as prototypical instances of disinformation. For instance, even if a political advertisement does not convince me that a particular candidate would dismantle Social Security, I might still not vote for her or him just because I am now uncertain about what she/he would do. Thus, we might want to amend my new analysis and simply require that disinformation be likely to create false beliefs *or* prevent true beliefs (and have the function of doing so).

APPLICATIONS OF THE ANALYSIS

Once we have a better understanding of exactly what disinformation is, we are in a much better position to deal with the serious threat to information quality that disinformation poses. My new analysis of disinformation can potentially help us to develop techniques for identifying disinformation and policies for deterring its spread.

Many information scientists (for example, Fallis & Frické, 2002; Kunst, Groot, Latthe, Latthe, & Khan, 2002) are searching for indicators of inaccurate information on the internet; that is, they are trying to identify features of websites that tend to be correlated with inaccuracy. Unfortunately, such research has consistently failed to differentiate among the various types of inaccurate information (for example, between intentionally misleading information and accidental falsehoods). This is a serious flaw, as the clues that suggest that someone is lying are unlikely to be the same clues that suggest that a person just does not know what she/he is talking about. Once we clarify what disinformation is and look specifically for indicators of it, we are likely to find better techniques for identifying it.

Since disseminating disinformation is still very close to lying in my new analysis, a lot of the vast research on lie detection can potentially be applied to disinformation detection. Researchers in lie detection have focused primarily on physiological indicators of deception, such as perspiration and high pulse rate (Vrij, 2008). However, we do not always come into direct physical contact with sources of disinformation; and even if we do, we are rarely in a position to give this source a polygraph test. But research is now being done to identify indicators of deception in both text (Newman, Pennebaker, Berry, & Richards, 2003; Rubin & Conroy 2012) and images (Farid, 2009, pp. 100–106). For example, researchers have used textual analysis to find that liars are somewhat less likely to use first-person pronouns.

In addition, this analysis of disinformation can be used create formal models of how disinformation spreads (Fallis, 2014; Skyrms, 2010, pp. 73–82; Sober, 1994, pp. 71–92; Tullock, 1967, pp. 133–143). One way in which disinformation can acquire the function of misleading is by the source systematically benefiting from its being misleading. In line with this, these formal models assume that whether a person will disseminate disinformation depends on the expected costs and benefits; more precisely, whether a person will disseminate disinformation depends on the costs of not being believed (weighted by the probability that this will happen), as compared with the benefits of being believed (weighted by the probability that this will happen). Such formal models can allow us to determine what sorts of things affect the amount of disinformation, and how they affect it. For instance, Fallis (2014) uses a formal model to show that the amount of disinformation tends to decrease when: a) the benefits to the recipient of believing what the source says when it is true decrease; b) the costs of believing what the source says when it is false increase; or c) the chances that the source has a motivation to deceive increase. Such results can potentially help us to devise policies that might deter people from disseminating disinformation.

CONCLUSION

Disinformation can cause significant harm if people are misled by it. But in order to address this critical threat to information quality, we first need to understand exactly what disinformation is. After surveying the various analyses that have been proposed, I have argued that disinformation is misleading information that has the function of misleading.

ACKNOWLEDGMENTS

For helpful feedback on earlier versions of this paper I would like to thank Allen Buchanan, Blaise Cronin, Andrew Dillon, Tony Doyle, Kenny Easwaran, Luciano Floridi, Martin Frické, Jonathan Furner, Richard Fyffe, Ken Herold, Birger Hjørland, Phyllis Illari, Cheryl Knott, Jens-Erik Mai, Neil Manson, Kay Mathiesen, Ken McAllister, Patti Montiel Overall, Bruce Schneier, Yotam Shmargad, David Livingstone Smith, Sille Obelitz Søre, Dan Zelinski, my graduate seminar on “Misinformation and Disinformation,” the anonymous referees for the iConference, and audiences at the School of Information Resources and Library Science at the University of Arizona (UA), the Philosophy Club at UA, a Workshop on Information Quality at the University of Hertfordshire, the 2009 Pacific Division Meeting of the American Philosophical Association, the 2009 iConference, the 2010 annual meeting of the American Society for Information Science and Technology, the 2013 Security and Human Behavior conference, and the 2014 iConference. This research was supported by a Research Professorship from the Social and Behavioral Sciences Research Institute at UA.

NOTES

1. A few information scientists (for example, Hernon, 1995, p. 134) take *misinformation* and *disinformation* to be mutually exclusive categories. However, it is more standard to use the term *misinformation* to refer to inaccurate information in general (Fox, 1983, p. 201), and to treat *disinformation* as a proper subset of misinformation (Floridi, 2011, p. 260; Skyrms, 2010, p. 80).
2. It is also important to be able to distinguish between the different types of disinformation, such as lies, spin, and even bullshit (Mearsheimer, 2011, p. 27).
3. *Conceptual analysis* has been criticized because of this reliance upon intuitions. For example, several people have argued that our intuitions about hypothetical cases cannot give us “a priori knowledge of necessary truths” (Melnyk, 2008, p. 267). Also, there is empirical evidence that such intuitions are fallible (Margolis & Laurence, 2011). But even if our intuitions can only provide us with fallible, a posteriori knowledge, they can still be used effectively to test proposed analyses of important concepts.
4. Admittedly, the term *disinformation* is relatively new compared with a term like *knowledge*. It is only about fifty years old. As a result, the meaning of *disinformation* may not be quite as fixed as the meaning of *knowledge*. Even so, we must have somewhat stable, shared intuitions about the use of the term; otherwise, we would not be able to effectively use it to communicate with one another (Jackson, 1998).
5. The fake thumbprint created in order to frame the “unfortunate” John Hector McFarlane for murder in one of the *Sherlock Holmes* stories would also count.
6. Along similar lines, Floridi (2011) points out that “a false friend” is “not a friend at all” (p. 183). But it is not clear how far we should rely upon such analogies in our conceptual analysis. Other examples cut the other way; for instance, false teeth *are* a kind of teeth.

Also, we are more likely to say “the information that you gave me was incorrect” than to say that it was not information.

7. *Disinformation* is my focus in this paper. Despite being false, jokes and sarcastic comments are probably not *misinformation* either. We should probably say that *misinformation* is information that is inaccurate and misleading; in fact, according to Skyrms (2010, p. 80), misinformation simply is misleading information.
8. This is analogous to a young child trying to poison a sibling by putting broccoli in his or her food. Assuming that he or she does not happen to be allergic to broccoli, the sibling is not actually put in danger despite the young child’s malicious intention.
9. Even if one agrees with Dretske and Floridi that information must be true, these examples show that some disinformation *is* information.
10. The information scientist Hernon (1995) performed a much earlier study that involved placing inaccurate information on the internet.
11. Birds and squirrels probably simply learn to associate making a certain vocalization and conspecifics running away and leaving food behind. Other species, such as fireflies, have *evolved* to send deceptive signals.
12. There is also a mechanism that reinforces the dissemination of the false claims made in the *Onion*. However, in this case, more readers are attracted to the website just because they find these claims amusing.
13. The source of such detrimental disinformation may not even be aware that the information is misleading. An extreme example is the racist statements made by the “white supremacist” Clayton Bigsby (played by comedian Dave Chappelle). Bigsby is black, but he is also blind; he does not know that he is black and, thus, that he is one of the targets of his own racism. (The other members of the KKK also do not know that he is black because they always wear hoods at meetings.)
14. This rough characterization will be sufficient for our purposes here. But a complete analysis of the concept of *disinformation* will require a careful analysis of the concept of *function*. Krohs and Kroes (2009) provide a selection of current philosophical research in this area.
15. People certainly treat the etiological functions of biological organisms (à la Millikan, 1984) and the design functions of artifacts (à la Dipert, 1993) as being two species of the same genus. As a result, several philosophers (for example, the various contributions to Krohs and Kroes’s *Functions in Biological and Artificial Worlds* [2009]) are looking to give a unified account of functions. But it is a difficult task to subsume one type of function under the other (see, for example, Vermaas & Houkes, 2003). Therefore, it may yet turn out that an analysis of disinformation must ultimately be disjunctive (that is, that there are two distinct ways in which information can be nonaccidentally misleading).
16. This is not necessarily an exhaustive list of the ways in which a piece of information can acquire the function of misleading. For instance, Bigsby’s racist statements (see note 13 above) are not intended to be misleading (since he believes what he is saying); also, he does not ultimately benefit from disseminating this misleading information (since he himself is black). Nevertheless, there is clearly a mechanism that reinforces the dissemination of this disinformation. The members of the KKK *do* benefit from its dissemination, and (not knowing that he is black) they give him all sorts of encouragement.
17. This is how many lie-detection techniques work (Vrij, 2008). Such techniques look for indications, such as nervousness or apprehension, of a “guilty mind.”
18. Thus, even if we bite the bullets and do not count them as disinformation per se, visual disinformation, true disinformation, side-effect disinformation, and adaptive disinformation should be included in our taxonomy of epistemically dangerous information.
19. Federal law on deceptive advertising tends to set a fairly low bar. For instance, as Carson (2010) notes, “ads that convey false claims to 15% or more of the target audience are usually found to be deceptive by Lanham courts” (p. 187).
20. I am assuming here that Moriarty does not intend to take advantage of some cognitive quirk that is unique to Holmes. In that case, it would be an accident if Watson is misled.

REFERENCES

- American Heritage. (2000). *The American Heritage dictionary of the English language* (4th ed.). Boston: Houghton Mifflin.

- American Library Association. (2005). Resolution on disinformation, media manipulation and the destruction of public information. Retrieved November 26, 2014, from <http://www.ala.org/aboutala/sites/ala.org.aboutala/files/content/governance/policymanual/updatedpolicymanual/ocrpdfofprpm/52-8disinformation.pdf>
- Austin, J. L. (1956). A plea for excuses. *Proceedings of the Aristotelian Society*, 57, 1–30.
- Bates, M. J. (2006). Fundamental forms of information. *Journal of the American Society for Information Science and Technology*, 57(3), 1033–1045.
- Buckland, M. K. (1991). Information as thing. *Journal of the American Society for Information Science* 42(5), 351–360.
- Calvert, P. J. (2001). Scholarly misconduct and misinformation on the World Wide Web. *Electronic Library*, 19(4), 232–240.
- Carson, T. L. (2010). *Lying and deception*. New York: Oxford University Press.
- Chisholm, R. M., & Feehan, T. D. (1977). The intent to deceive. *Journal of Philosophy*, 74(3), 143–159.
- Dipert, R. R. (1993). *Artifacts, art works, and agency*. Philadelphia: Temple University Press.
- Dretske, F. I. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.
- Dretske, F. I. (1983). Précis of *Knowledge and the flow of information*. *Behavioral and Brain Sciences*, 6(1), 55–90.
- Fallis, D. (2009). A conceptual analysis of disinformation. In *iConference Proceedings*. Retrieved November 26, 2014, from <http://hdl.handle.net/2142/15205>
- Fallis, D. (2011). Floridi on disinformation. *Etica & Politica*, 13(2), 201–214.
- Fallis, D. (2014). Epistemic values and disinformation. In A. Fairweather (Ed.), *Virtue epistemology naturalized: Bridges between virtue epistemology and philosophy of science* (pp. 159–179). Dordrecht: Springer.
- Fallis, D., & Frické, M. (2002). Indicators of accuracy of consumer health information on the internet. *Journal of the American Medical Informatics Association*, 9(1), 73–79.
- Fallon, K. (2012). Fooled by “The Onion”: 9 most embarrassing fails. *Daily Beast*. Retrieved November 26, 2014, from <http://www.thedailybeast.com/articles/2012/09/29/fooled-by-the-onion-8-most-embarrassing-fails.html>
- Farid, H. (2009). Digital doctoring: Can we trust photographs? In B. Harrington (Ed.), *Deception: From ancient empires to internet dating* (pp. 95–108). Palo Alto, CA: Stanford University Press.
- Farquhar, M. (2005). *A treasury of deception*. New York: Penguin.
- Feldman, R. (2003). *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Fetzer, J. H. (2004). Disinformation: The use of false information. *Minds and Machines*, 2(14), 231–240.
- Fiore, F., & Francois, J. (2002, September 6). Unwitting collaborators, part 12: Disinformation changing web site contents. *Informit*. Retrieved November 26, 2014, from <http://www.informit.com/articles/article.aspx?p=29255>
- Fitzpatrick, S. (2009). The primate mindreading controversy: A case study in simplicity and methodology in animal psychology. In R. Lurz (Ed.), *The philosophy of animal minds* (pp. 258–277). Cambridge: Cambridge University Press.
- Floridi, L. (1996). Brave.net.world: The internet as a disinformation superhighway? *Electronic Library*, 14, 509–514.
- Floridi, L. (2005). Semantic conceptions of information. In *Stanford Encyclopedia of Philosophy*. Retrieved November 26, 2014, from <http://plato.stanford.edu/entries/information-semantic>
- Floridi, L. (2011). *The philosophy of information*. New York: Oxford University Press.
- Floridi, L. (2012). Steps forward in the philosophy of information. *Etica & Politica*, 14(1), 304–310.
- Fowler, B., Franklin, C., & Hyde, R. (2001, February 28). Internet Securities Fraud: Old Trick, New Medium. *Duke Law & Technology Review*. <http://dltr.law.duke.edu/2001/02/28/internet-securities-fraud-old-trick-new-medium/>
- Fox, C. J. (1983). *Information and misinformation: An investigation of the notions of information, misinformation, informing, and misinforming*. Westport, CT: Greenwood Press.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23: 121–123.
- Graham, P. (2010). Testimonial entitlement and the function of comprehension. In A. Hadlock, A. Miller, & D. Pritchard (Eds.), *Social epistemology* (pp. 148–174). New York: Oxford University Press.

- Halavais, A. (2004). The Isuzu experiment. Retrieved November 26, 2014, from <http://alex.halavais.net/the-isuzu-experiment>
- Hancock, J. T. (2007). Digital deception: When, where and how people lie online. In A. Joinson, K. McKenna, T. Postmes, and U.-D. Reips (Eds.), *Oxford handbook of internet psychology* (pp. 289–301). Oxford: Oxford University Press.
- Hernon, P. (1995). Disinformation and misinformation through the internet: Findings of an exploratory study. *Government Information Quarterly*, 12(2), 133–139.
- Jackson, F. (1998). *From metaphysics to ethics*. New York: Oxford University Press.
- Karlova, N. A., & Fisher, K. E. (2013). A social diffusion model of misinformation and disinformation for understanding human information behavior. *Information Research*, 18(1). Retrieved November 26, 2014, from <http://informationr.net/ir/18-1/paper573.html>
- Kingsbury, J., & McKeown-Green, J. (2009). Definitions: Does disjunction mean dysfunction? *Journal of Philosophy*, 106(10), 568–585.
- Krohs, U., & Kroes, P. (Eds.). (2009). *Functions in biological and artificial worlds*. Cambridge, MA: MIT Press.
- Kunst, H., Groot, D., Latthe, P. M., Latthe, M., & Khan, K. S. (2002). Accuracy of information on apparently credible websites: Survey of five common health topics. *British Medical Journal*, 324(7337), 581–582.
- Lynch, C. A. (2001). When documents deceive: Trust and provenance as new factors for information retrieval in a tangled web. *Journal of the American Society for Information Science and Technology*, 52(1), 12–17.
- Mahon, J. E. (2008). The definition of lying and deception. In *Stanford Encyclopedia of Philosophy*. Retrieved November 26, 2014, from <http://plato.stanford.edu/entries/lying-definition>
- Margolis, E., & Laurence, S. (2011). Concepts. In *Stanford Encyclopedia of Philosophy*. Retrieved November 26, 2014, from <http://plato.stanford.edu/entries/concepts>
- Mearsheimer, J. J. (2011). *Why leaders lie*. New York: Oxford University Press.
- Melnyk, A. (2008). Conceptual and linguistic analysis: A two-step program. *Noûs*, 42(2), 267–291.
- Millikan, R. G. (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.
- Monmonier, M. (1991). *How to lie with maps*. Chicago: University of Chicago Press.
- Moran, R. (2005). Getting told and being believed. *Philosophers' Imprint*, 5(5). Retrieved November 26, 2014, from <http://www.philosophersimprint.org/005005>
- Nesta, F., & Blanke, H. (1991). Warning: Propaganda! *Library Journal*, 116(9), 41–43.
- Newman, M. L., Pennebaker, J. W., Berry, D. S., & Richards, J. M. (2003). Lying words: Predicting deception from linguistic styles. *Personality and Social Psychology Bulletin*, 29(5), 665–675.
- Piper, P. S. (2002). Web hoaxes, counterfeit sites, and other spurious information on the internet. In A. P. Mintz (Ed.), *Web of deception* (pp. 1–22). Medford, NJ: Information Today.
- Pratchett, T., Stewart, I., & Cohen, J. (1999). *The science of Discworld*. London: Ebury Press.
- Reddy, S. (2013, February 18). “I don’t smoke, doc,” and other patient lies. *Wall Street Journal*. Retrieved November 26, 2014, from <http://online.wsj.com/article/SB10001424127887323478004578306510461212692.html>
- Rowe, M. W. (2012). The problem of perfect fakes. *Royal Institute of Philosophy Supplement*, 71, 151–175.
- Rubin V. L., & Conroy, N. (2012). Discerning truth from deception: Human judgments and automation efforts. *First Monday*, 17(3). Retrieved November 26, 2014, from <http://firstmonday.org/ojs/index.php/fm/article/view/3933/3170>
- Scarantino, A., & Piccinini, G. (2010). Information without truth. *Metaphilosophy*, 41(3), 313–330.
- Schauer, F., & Zeckhauser, R. (2009). Paltering. In B. Harrington (Ed.), *Deception: From ancient empires to internet dating* (pp. 38–54). Palo Alto, CA: Stanford University Press.
- Skinner, S., & Martin, B. (2000). Racist disinformation on the World Wide Web: Initial implications for the LIS community. *The Australian Library Journal*, 49(3), 259–269.
- Skyrms, B. (2010). *Signals: Evolution, Learning, and Information*. Oxford: Oxford University Press.
- Sober, E. (1994). *From a biological point of view: Essays in evolutionary philosophy*. New York: Cambridge University Press.
- Sowards, S. W. (1988). Historical fabrications in library collections. *Collection Management*, 10(3–4), 81–88.
- Tulloch, G. (1967). *Toward a mathematics of politics*. Ann Arbor: University of Michigan Press.

- Twain, M. (2002). My first lie, and how I got out of it. In *The man that corrupted Hadleyburg* (pp. 159–170). New York: Prometheus Books. (Original work published 1900)
- Vermaas, P. E., & Houkes, W. (2003). Ascribing functions to technical artefacts: A challenge to etiological accounts of functions. *British Journal for the Philosophy of Science*, 54, 261–289.
- Vrij, A. (2008). *Detecting lies and deceit: Pitfalls and opportunities* (2nd ed.). Chichester, UK: Wiley.
- Walsh, J. (2010). Librarians and controlling disinformation: Is multi-literacy instruction the answer? *Library Review*, 59(7), 498–511.
- Weatherson, B. (2003). What good are counterexamples? *Philosophical Studies*, 115(1), 1–31.
- Whitty, M. T., Buchanan, T., Joinson, A. N., & Meredith, A. (2012). Not all lies are spontaneous: An examination of deception across different modes of communication. *Journal of the American Society for Information Science and Technology*, 63(1), 208–216.

Don Fallis is a professor of information resources and adjunct professor of philosophy at the University of Arizona. His research areas are epistemology, philosophy of information, and philosophy of mathematics. His articles have appeared in *Library Trends*, *The Information Society*, *Journal of the American Society for Information Science and Technology*, *Philosophical Studies*, *American Mathematical Monthly*, and *British Journal for the Philosophy of Science*, and he has a contribution titled “Wikipistemology” in *Social Epistemology: Essential Readings*. His recent articles on lying and deception include “What is Lying?” in the *Journal of Philosophy*. He has also discussed lying and deception on *Philosophy TV* and in several volumes of the *Popular Culture and Philosophy* series.