
On Verifying the Accuracy of Information: Philosophical Perspectives

DON FALLIS

The sooner a man begins to verify all he hears
the better it is for him.

—George Gurdjieff

Trust but verify.

—Ronald Reagan

ABSTRACT

HOW CAN ONE VERIFY THE ACCURACY OF RECORDED INFORMATION (e.g., information found in books, newspapers, and on Web sites)? In this paper, I argue that work in the epistemology of testimony (especially that of philosophers David Hume and Alvin Goldman) can help with this important practical problem in library and information science. This work suggests that there are four important areas to consider when verifying the accuracy of information: (i) authority, (ii) independent corroboration, (iii) plausibility and support, and (iv) presentation. I show how philosophical research in these areas can improve how information professionals go about teaching people how to evaluate information. Finally, I discuss several further techniques that information professionals can and should use to make it easier for people to verify the accuracy of information.

PHILOSOPHY OF INFORMATION (PI)

PI is “the philosophical field concerned with the critical investigation of the conceptual nature and basic principles of information, including its dynamics, utilisation, and sciences” (Floridi, 2002b, p. 123).¹ Luciano Floridi (2002a) and Ken Herold (2001) recently looked at the broad connec-

Don Fallis, School of Information Resources and Library Science, University of Arizona, 1515 East First Street, Tucson, AZ 85719

LIBRARY TRENDS, Vol. 52, No. 3, Winter 2004, pp. 463–487

© 2004 The Board of Trustees, University of Illinois

tions between library and information science (LIS) and PI. For example, Floridi (2002a, pp. 47–48) considers whether PI might serve as the theoretical foundation for which library and information scientists have long been searching. Of course, it is somewhat rare for philosophers to explicitly address LIS issues.² Also, PI has only recently been identified as a distinct field of inquiry. Even so, philosophers have been working on many issues of concern to LIS for centuries (cf. Floridi, 2002a, p. 44).

This paper focuses on how PI can help with one specific practical concern for LIS, namely, how one can verify the accuracy of recorded information. In other words, how one can determine if the information found in a book, in a newspaper, on a Web site, etc. is accurate.

Notably, library and information scientists typically talk about evaluating the quality of information rather than about verifying the accuracy of information (see, e.g., Alexander & Tate, 1999; Cooke, 1999). In fact, accuracy is usually just one of the traditional criteria for evaluating the quality of information (see, e.g., Alexander & Tate, 1999, pp. 11–13; Cooke 1999, pp. 60–62).

Even for library and information scientists, however, accuracy is the *sine qua non* of quality information sources. For example, consider some of the other traditional criteria for evaluating the quality of information. The main reason that we are interested in finding information sources that are authoritative, objective, and current is that we think that they are more likely to be accurate. In other words, these criteria are indicative of information quality precisely because they are indicative of information accuracy.

Library and information scientists are legitimately concerned with quality issues that do go beyond accuracy, such as the accessibility, relevance, comprehensibility, and navigability of information sources. However, as Peter Hernon (1995, p. 133) points out, “it is not enough that information is readily available; before relying on any data or information, it may be important to ascertain, for example, the veracity of the content.”

In this paper, I appeal mainly to the work of David Hume (1748/1977) and Alvin Goldman (1999 and 2001) on the epistemology of testimony (as well as to some work in game theory).³ Their work suggests four important areas to consider when verifying the accuracy of information: (i) authority, (ii) independent corroboration, (iii) plausibility and support, and (iv) presentation. I show how work in these areas can improve how information professionals go about teaching people how to evaluate information. In addition, I argue that information professionals can and should use some further important techniques to make it easier for people to verify the accuracy of information.

THE PROBLEM OF INACCURATE INFORMATION

Even fairly reliable information sources contain some amount of inaccurate information. Famously, the *Chicago Tribune* mistakenly reported that

Dewey had defeated Truman in the 1948 presidential campaign. More recently, Dan Rather mistakenly reported that James Brady had died after being shot by John Hinkley (cf. Frické 1997, p. 887). However, since almost anyone can post almost anything on the Internet with no editorial control, we might expect much more inaccurate information on the Internet.⁴ In fact, empirical studies have found a considerable amount of inaccurate information on the Internet (see, e.g., Impicciatore et al., 1997; Connell & Tipple, 1999; Berland et al., 2001).⁵

The reason that inaccurate information is a problem is that people can often be misled by it. And the risks here are not just epistemic. People use the information that they have to make practical decisions. If people are misled by inaccurate information, it can cause serious harm to their finances (cf. Fowler et al., 2001) and their health (cf. Kiley, 2002). In addition, while some people may be too credulous, other people may be too skeptical. Because they are worried about being misled, some people may fail to believe accurate information that it would have been beneficial for them to believe.

Of course, the mere fact that an information source contains some amount of inaccurate information is not necessarily a problem (cf. Wachbroit, 2000, p. 11). As long as people can distinguish the accurate from the inaccurate information, they will not be misled. Unfortunately, it often can be very difficult for people to identify inaccurate information (cf. Cerf, 2002). For example, with the latest Web development software, almost anyone can publish very professional-looking Web sites. As Silberg et al. (1997, p. 1244) point out, the Internet "is a medium in which anyone with a computer can serve simultaneously as author, editor, and publisher and can fill any or all of these roles anonymously if he or she so chooses. In such an environment, novices and savvy Internet users alike can have trouble distinguishing the wheat from the chaff, the useful from the harmful."

Thus, people need to be able to distinguish the accurate information from the inaccurate information. In other words, they need to be able to verify the accuracy of information. Since the problem of inaccurate information seems to be most pressing on the Internet, this paper focuses specifically on how to verify the accuracy of information on the Internet. Even so, almost all of the points that are made can be applied generally to verifying the accuracy of information from any source.

Library and information scientists have responded to the problem of inaccurate information on the Internet primarily by publishing guidelines for evaluating information (see, e.g., Ambre et al., 1997; Silberg et al., 1997; Wilkinson et al., 1997; Alexander & Tate, 1999; Cooke, 1999; Smith, A., 2002). These guidelines provide people with a list of features of Web sites that are supposed to be indicators of accuracy (e.g., the author is identified, the author is an authority on the topic, no advertising appears, no spelling or grammatical errors are present, the Web site is up-to-date, authoritative

references are cited). As I discuss in the following sections, work on the epistemology of testimony provides a conceptual framework that encompasses such efforts to deal with the problem of inaccurate information. This conceptual framework explains why these guidelines work and suggests how they can be improved. In addition, having such a conceptual framework can make it easier to apply and to communicate these guidelines.

Finally, notably, the problem of inaccurate information is not *sui generis*. This problem is analogous to a number of other problems that people confront. For example, to deal effectively with the problem of counterfeit currency, people need to be able to distinguish authentic currency from counterfeit currency (cf. Bureau of Engraving and Printing, 2002). Similarly, to deal effectively with the problem of low-quality products, consumers need to be able to distinguish high-quality products from low-quality products (cf. Baird et al., 1994, pp. 122–125). I explain how strategies for dealing with these other problems have the potential to be heuristically valuable when we confront the problem of inaccurate information.

THE EPISTEMOLOGY OF TESTIMONY

The problem of how to verify the accuracy of recorded information is a special case of the problem of how to verify the accuracy of testimony. For example, suppose that we want to know the height of the Eiffel Tower. Very few people have the resources and expertise to measure the Eiffel Tower for themselves. The rest of us have to look this information up in a book or on a Web site. But when we get this information from a book or a Web site, we are relying on the testimony of the author.

The epistemology of testimony is important because a large amount of the information that we have about the world comes from others rather than from direct observation (cf. Lipton, 1998, p. 2). As Hume (1748/1977, p. 74) puts it, “there is no species of reasoning more common, more useful, and even necessary to human life, than that which is derived from the testimony of men, and the reports of eye-witnesses and spectators.” And, in particular, a large amount of the information that we get from others comes to us as recorded information (e.g., in books, in newspapers, and on Web sites).

While the epistemology of testimony certainly has not been the main focus of traditional epistemology (cf. Goldman, 1999, p. 4), a lot of work has been done in this area (see, e.g., Giedymin, 1963; Hardwig, 1985; Coady, 1992; Lipton, 1998; Goldman, 1999, pp. 103–130; Goldman, 2001;). In fact, work in this area goes back to the ancient Greeks. Plato (380 BC/2002, p. 170d), for example, asked whether one can “examine another man’s claim to some knowledge, and make out whether he knows or does not know what he says he knows.”

One of the most influential early discussions of the epistemology of testimony is in David Hume’s *An Enquiry Concerning Human Understanding* (1977), originally published in 1748. In the chapter “Of Miracles,” Hume

develops a general framework that can be used to evaluate any kind of testimony. In particular, this framework can be applied to verifying the accuracy of recorded information.⁶

According to Hume (1748/1977, p. 77), “when any one tells me, that he saw a dead man restored to life, I immediately consider with myself, whether it be more probable, that this person should either deceive or be deceived, or that the fact, which he relates, should really have happened.” Hume is explicitly concerned here with evaluating testimony that a miracle has occurred, but his strategy can just as easily be used to evaluate, for example, the claim on a Web site that a particular treatment for a child with a fever is safe and effective. In both cases, we should consider all of the available evidence (including the fact of the testimony itself) and determine what the best explanation of this evidence is. This strategy is known as “inference to the best explanation.” Among philosophers, it is a common way to analyze scientific inference, but it also can be applied to the evaluation of testimony (cf. Lipton, 1998, p. 27). For example, what is the best explanation of the fact that the Web site is claiming that this treatment is safe and effective? Is it more likely that the Web site is promoting an ineffective treatment either because its author is not medically qualified or because the sponsor is selling pharmaceutical products? Alternatively, is it more likely that the treatment really is safe and effective?

In addition to inferring the best explanation, Hume (1748/1977, p. 73) claims that “a wise man . . . proportions his belief to the evidence.”⁷ In general, if the evidence that a claim is accurate is greater than the evidence that the claim is inaccurate, then we should be inclined to think that the claim is accurate (cf. Hume, 1748/1977, pp. 73–74).⁸ However, our degree of confidence in the accuracy of the claim should depend on how much the evidence of accuracy exceeds the evidence of inaccuracy. As Hume (1748/1977, p. 74) puts it, “a hundred instances or experiments on one side, and fifty on another, afford a doubtful expectation of any event; though a hundred uniform experiments, with only one contradictory, reasonably beget a pretty strong degree of assurance.”

But what evidence should we consider when we are trying to verify the accuracy of a piece of information?⁹ Hume has several suggestions as to what evidence is relevant. For example, Hume (1748/1977, p. 75) says that “we entertain a suspicion concerning any matter of fact, when the witnesses contradict each other; when they are but few, or of a doubtful character; when they have an interest in what they affirm; when they deliver their testimony with hesitation, or on the contrary, with too violent asseverations” (cf. Locke, 1690/1975, pp. 662–663). In fact, his suggestions sound a lot like the published guidelines for evaluating information on the Internet.

The evidence that Hume suggests that people consider can be divided into roughly four categories. These categories will be examined in more detail in the following sections. However, before we move on to these de-

tails, two important points should be emphasized about Hume's recommendations for evaluating testimony.

First, Hume is engaged in a normative project. That is, he is making recommendations for how people ought to go about verifying the accuracy of information. The published guidelines for evaluating information on the Internet are also engaged in this same normative project.

By contrast, a number of recent LIS articles (see, e.g., Fogg et al., 2001; Eysenbach & Köhler, 2002; Rieh, 2002; Wathen & Burkell, 2002) have looked at how people actually do make judgments about information quality. Of course, the features of Web sites that people actually find credible need not coincide with the features that they ought to find credible. In fact, determining what people find credible is most directly relevant to helping authors devise ways to convince their audience whether or not what they are saying is accurate.

Second, the goal of following Hume's recommendations is that we end up with true beliefs. In other words, the goal is to acquire beliefs that correspond to reality.¹⁰ Philosophers typically take this to be the goal of information seekers (see, e.g., Goldman, 1999, p. 3).¹¹ John Locke (1690/1975, p. 697), for example, explicitly states that the reason that we should proportion our belief to the evidence is so that we will end up with true beliefs.

Library and information scientists, however, are much less likely to take this to be the goal of information seekers (cf. Fallis, 2000, p. 314). Jesse Shera (1970, p. 97), for example, says that "*false* knowledge . . . is still knowledge, it is knowable and known." However, as a number of library and information scientists have recently argued (see, e.g., Frické, 1997, p. 887; Meola, 2000, p. 174; Doyle, 2001, pp. 62–63;), information seekers often do have the goal of acquiring true beliefs. For example, a student writing a report on the Eiffel Tower wants to know how tall the Eiffel Tower really is. In other words, she is after the truth. Similarly, a parent wants to know whether a particular treatment for a child with a fever really is safe and effective. In fact, it does not really make sense for someone to bother about verifying the accuracy of information unless acquiring true beliefs is her goal.¹²

Of course, this is not to say that information seekers exclusively seek true beliefs (cf. Goldman, 1999, p. 26). Walter Kaufmann (1977, pp. 47–83), for example, argues that the goal of reading should be to acquire more than just true beliefs (e.g., understanding).¹³ But even if it is not our ultimate goal, acquiring true beliefs is often a critical means to this further end (cf. Bruce, 2000, p. 109).

Finally, even if our goal is simply to acquire true beliefs, it is important to note that we are not just interested in the accuracy of the information. A person can be misled by incomplete information as well as by inaccurate information. Thus, we also are interested in the completeness of the information (cf. Frické, 1997; Fallis & Frické, 2002, pp. 74–75).¹⁴ This paper, however, focuses specifically on how to verify the accuracy of information.

WHO TESTIFIES: AUTHORITY

Library and information scientists emphasize that people should consider the source of a piece of information when trying to verify its accuracy. In particular, people are advised to determine the authority of the information source (see, e.g., Wilson, 1983; Alexander & Tate, 1999, p. 11; Cooke, 1999, pp. 58–60). This is what Hume (1748/1977, p. 75) had in mind when he said that we should consider the “character . . . of the witnesses.”

Deciding what to believe based on who said it is referred to by philosophers as an appeal to authority. An appeal to authority is often listed in introductory logic texts as a fallacy. However, while some appeals to authority are fallacious (e.g., taking medical advice from someone who plays a doctor on TV), others are legitimate (cf. Salmon, 1995, p. 105; Goldman, 2001, p. 88).¹⁵ An appeal to authority is legitimate whenever our source is likely to be providing us with accurate information (i.e., if she is a reliable source on the topic in question).

As Jerzy Giedymin (1963, pp. 288–289) points out, there are essentially two ways to determine whether an information source is reliable. First, has this information source usually provided accurate information in the past? If an individual has been right about things in the past, then she is more likely to be right about things now (cf. Hume, 1748/1977, p. 73). Second, does anything suggest that this information source would not provide accurate information in this particular case? For example, the published guidelines on evaluating information typically advise people to determine if the information source has an obvious bias (see, e.g., Ambre et al., 1997; Alexander & Tate, 1999, p. 13).¹⁶

Regarding “past track record” (Goldman, 2001, p. 106), many circumstances exist where people can directly determine whether an information source has provided accurate information in the past (cf. Goldman, 2001, pp. 106–108). For example, while a student may not know how tall the Eiffel Tower is, she will probably know that the Eiffel Tower is in Paris. Thus, if a Web site incorrectly claims that the Eiffel Tower is in Brussels, she probably should not trust this Web site when it claims that the Eiffel Tower is exactly 300 meters high.

However, a couple of practical difficulties arise with trying to directly determine the past track record of an information source. For one thing, people may not have had a sufficient number of previous interactions with this particular information source (cf. Lipton, 1998, p. 15). In fact, given the size of the Internet, people often will be consulting particular Web sites for the first time. Also, people may not have the expertise to judge the accuracy of the information previously provided by this information source (cf. Goldman, 2001, p. 106). As a result, if people only rely on their own past experience with information sources, they will only be able to verify the accuracy of a limited amount of the information on the Internet.

Fortunately, people also can rely on the experience that other people

have had with an information source. In other words, people can use the testimony of others to determine the reliability of an information source (cf. Goldman, 2001, p. 97). For example, the published guidelines on evaluating information typically advise people to determine if the information source has a reputation for reliability (see, e.g., Wilkinson et al., 1997, p. 55; Cooke, 1999, p. 61). Based on sources' reputations, someone seeking financial information can be reasonably confident in the accuracy of information found on the *Wall Street Journal* Web site or the Bloomberg Web site.

In fact, even if an information source does not have a well-established reputation, other factors can determine if people recommend or endorse this information source.¹⁷ For example, the published guidelines on evaluating information typically advise people to determine if the information source has appropriate credentials (see, e.g., Ambre et al., 1997). As Goldman (2001, p. 97) points out, conferring credentials, such as academic degrees and professional accreditations, is a way in which experts commonly endorse other experts. In addition, people can determine whether an information source on the Internet is endorsed by others by looking at how many Web sites link to this Web site (cf. Burbules, 2001, p. 444). This is analogous to using citation counts as an indication of the scholarly quality of a journal (cf. Lee et al., 2002).

Bias is just one of a number of features that might suggest that an information source would not provide accurate information in a particular case. As Hume (1748/1977, p. 77) suggests when he says that a witness may "either deceive or be deceived," these features fall roughly into two categories. First, is there any indication that the witness was not sincere in her testimony?¹⁸ Second, is there any indication that the witness was not in a position to know the fact that she is testifying to? This is a standard distinction made by philosophers.¹⁹ It is basically the same distinction that information scientists have in mind when they talk about disinformation and misinformation (see, e.g., Herson, 1995, p. 134).

Finally, everything that has been discussed so far in this section (e.g., past track record, bias, reputation) is clearly an epistemically relevant consideration. However, philosophers have claimed that some unexpected features also might actually help people determine whether an information source is reliable. For example, Ashley McDowell (2002, pp. 60–61) argues that the general trustworthiness of an individual (e.g., she keeps her promises) may be indicative of epistemic trustworthiness. Linda Alcoff (1999) even argues that the social identity (e.g., culture, race, gender) of an individual may have an effect on her epistemic credibility.

HOW MANY TESTIFY: INDEPENDENT CORROBORATION

In the last section, we looked at ways to determine the reliability of a single source of information. However, in addition to the character of the witness, Hume (1748/1977, p. 75) also notes that people should pay atten-

tion to the “number of the witnesses.” This is because it is much more likely that one individual will “deceive or be deceived” than that several individuals will “deceive or be deceived” in exactly the same way. For this reason, several philosophers have noted that the agreement of a number of experts on a topic can be an important indicator of accuracy (see, e.g., Goldman, 1987, p. 122; Salmon, 1995, p. 105). This suggests that another technique for verifying the accuracy of a piece of information is to see if other information sources corroborate the original source of the information (cf. Burbules, 2001, p. 446; Wilkinson et al., 1997, p. 56).²⁰

Notably, however, agreement between information sources is not always an indication that their information is accurate. It depends on how these different sources got their information. In particular, if they all got their information from the same place, then ten sources saying the same thing is no better evidence than one source saying it. Goldman (2001, p. 101) refers to such information sources as “nondiscriminating reflectors.” In a similar vein, according to Wachbroit (2000, p. 13), “Wittgenstein in *On Certainty* presents the [absurd] image of someone trying to check a story in a newspaper by buying other copies of the same newspaper and reading the story again.”

This issue turns out to be especially important on the Internet since it is so easy for the very same information to be copied by several different Web sites. For example, in one of the Internet scams reported by Fowler et al. (2001), the news service Internet Wire posted what turned out to be a fraudulent press release about the Emulex Corporation. The inaccurate information in this press release was then quickly picked up by several other news services, such as Bloomberg, CBS Marketwatch, and Dow Jones. As a result, those investors that did try to verify the accuracy of the information by checking multiple sources still ended up being misled.

The same sort of replication occurs with consumer health information. For example, over thirty different Web sites (including the National Library of Medicine, 2002) provide the same information word for word on how to treat children with fever. In this case, the information seems to be accurate. However, the fact that all of these sites corroborate each other still does nothing to help us verify that the information is accurate. Agreement between sources should not increase our degree of confidence in the accuracy of a piece of information unless those sources are independent (cf. Giedymin, 1963, p. 291).²¹

Of course, information sources do not always agree with each other. In fact, it is fairly easy to find conflicting information (e.g., about the height of the Eiffel Tower) from different sources on the Internet (cf. Burbules, 2001, p. 452). If sources do conflict, then people simply have to determine which source is more reliable (or use some of the other techniques for verifying the accuracy of the information).

WHAT THEY TESTIFY TO: PLAUSIBILITY AND SUPPORT

Library and information scientists emphasize that people should consider who the source of the information is when trying to verify its accuracy. Philosophers, however, are much more likely to emphasize that people should look at what the information is. In particular, people are advised to consider the plausibility of a claim and the reasons offered in support of the claim.²²

In addition to knowing about the reliability of the source, an information seeker typically knows many other things about the world that should be taken into account. If it is very unlikely that a piece of information is accurate given everything else that we know about the topic in question, then we should be inclined to think that the information is inaccurate (cf. Lipton, 1998, p. 25). As Hume (1748/1977, p. 75) puts it, "the evidence, resulting from the testimony, admits of a diminution, greater or less, in proportion as the fact is more or less unusual" (cf. Locke 1690/1975, p. 663).²³ So, for example, a Web site that depicts a city in Minnesota "as a tropical paradise" (Piper, 2000) is immediately suspect.

In addition to considering the plausibility of the claim, philosophers emphasize that people need to consider the reasons (if any) that an information source offers in support of the claim (see, e.g., Goldman, 1999, pp. 130–160; Goldman, 2001, pp. 93–94).²⁴ The reasons offered often can provide the best evidence for the accuracy of the claim. This is certainly how things are supposed to work in science and mathematics, for example. If a mathematician wants to determine whether a mathematical claim is true, she is not going to start by checking the credentials of the person making the claim. She is going to check the proof that this person has given.

Under ideal circumstances, this is the sort of evidence that people would use to verify the accuracy of information. However, once again, people will not always have sufficient expertise to evaluate the plausibility of a claim or the reasons offered in support of the claim (cf. Goldman, 2001, p. 94).²⁵ This may explain why library and information scientists have not focused on this particular technique for verifying the accuracy of information. In those cases where people lack sufficient expertise, they will have to fall back on more indirect evidence of accuracy or inaccuracy (e.g., considering who testifies and how they testify).

HOW THEY TESTIFY: PRESENTATION

In addition to who testifies and what they testify to, Hume (1748/1977, p. 75) notes that people should pay attention to the "manner of their delivering their testimony." How a witness testifies is often indicative of the reliability of this witness (cf. section 4 above).²⁶

Numerous features of Web sites proposed as indicators of accuracy fall into this category. For example, does the Web site engage in any advertis-

ing? Advertising may indicate a lack of objectivity (cf. Alexander & Tate, 1999, p. 27). The idea is that the desire to sell products might override the desire to tell the truth. Also, does the Web site contain any spelling or grammatical errors? Such mistakes may indicate a lack of concern for quality and accuracy (cf. Wilkinson et al., 1997, p. 57; Cooke, 1999, p. 61): The idea is that if someone is careful enough to get the spelling right, she is more likely to be careful enough to get the facts right. Finally, does the Web site cite authoritative references (see, e.g., Ambre et al., 1997; Cooke, 1999, p. 61)?

These are just a few of the features of Web sites that have been proposed as indicators of accuracy in the published guidelines for evaluating information. Not enough space is available here to give an exhaustive list. However, it is possible to identify three general constraints that indicators of accuracy must satisfy.²⁷

First, an indicator of accuracy clearly must be correlated with information being accurate. In other words, a Web site that displays an indicator of accuracy (e.g., lack of advertising) must be more likely to contain accurate information than a Web site that does not display that indicator (cf. Fallis, 2000, p. 307). For example, Fallis and Frické (2002, p. 76) found that a Web site with accurate health information was over three times more likely to display the Health on the Net Foundation's (2002a) HONcode logo than a Web site with inaccurate health information.

Unfortunately, it is not clear that the features of Web sites that are usually proposed as indicators of accuracy are really correlated with accuracy. For one thing, most of the published guidelines for evaluating information are not based on empirical data. Instead, these guidelines are often based on surveys of "Web experts" (see, e.g., Ambre et al., 1997; Wilkinson et al., 1997). However, it is not clear that the features of Web sites that these "Web experts" believe to be indicators of accuracy really are indicators of accuracy. Even worse, in at least some cases, these guidelines have "clearly drawn from one another" (cf. Burbules, 2001, p. 445). In other words, they are nondiscriminating reflectors (cf. section 5 above).

An empiricist such as Hume, however, would certainly insist on basing such guidelines on empirical data. For example, with respect to testimony in general, Hume (1748/1977, p. 75) says that "the reason, why we place any credit in witnesses and historians, is not derived from any connexion, which we perceive a priori, between testimony and reality, but because we are accustomed to find a conformity between them." Unfortunately, the few studies (see, e.g., Griffiths & Christensen, 2000; Fallis & Frické, 2002; Kunst et al., 2002) that have empirically tested the features of Web sites that are usually proposed as indicators of accuracy have not found them to be correlated with accuracy. According to Kathleen Griffiths and Helen Christensen (2000, p. 1515), "currently popular criteria for evaluating the quality of Web sites were not indicators of content quality."

Second, people must put the right amount of faith in an indicator of

accuracy (cf. Hume, 1748/1977, p. 75). In other words, they need to proportion their belief to the evidence. If people do not put the right amount of faith in an indicator of accuracy, they can still end up being too credulous (or too skeptical).

Indicators of accuracy are rarely guarantees of accuracy (cf. Hume, 1748/1977, p. 74). As a result, there is usually plenty of room for people to overestimate the degree to which an indicator of accuracy is actually correlated with accuracy. Unfortunately, the published guidelines for evaluating information do not tell people how much faith to put in the features of Web sites that are usually proposed as indicators of accuracy. Of course, since their recommendations are not based on empirical data, it is not clear how they could do so.²⁸

Third, an indicator of accuracy must be correlated with accuracy in a "robust" way. In particular, it must be difficult for the author of a Web site to "fake" an indicator of accuracy.

James Perry (1985, p. 251) once made the tongue-in-cheek suggestion that "the presence of a colon in the title of a paper is the primary correlate of scholarship." While titular colonicity is almost certainly not correlated with accuracy, it does provide a nice example of something that is definitely not correlated with accuracy in a "robust" way. This is because it is very easy for an author to put a colon in the title of her paper to gain credibility even if her information is inaccurate. As Perry (1986, p. 177) notes, "when an evaluation technique has been developed and tested, it will become familiar to the evaluatees as well as the evaluators. Thus, its utility will decrease as evaluatees change their practices to gain higher evaluations" (cf. Burbules, 2001, p. 445). In other words, even if titular colonicity were highly correlated with information being accurate, it would probably not remain correlated for long.

The same sort of worry arises with respect to indicators of accuracy on the Internet. Many Web sites containing inaccurate information might simply be following the suggestions from Fogg et al. (2001) on "what makes Web sites credible," for example. In fact, it should be noted that some Web sites containing inaccurate information do not simply try to look reputable. They actually try to look like somebody else. For example, in one of the Internet scams reported by Fowler et al. (2001), a Web site was designed to look like the Bloomberg Web site to fool investors. Interestingly enough, such Web sites are commonly referred to as "counterfeit sites" (cf. Piper, 2000).

Even so, it is possible for indicators to be correlated with accuracy in a "robust" way (cf. Fallis & Frické, 2002, p. 78). The author of a Web site needs to do something that she would be unable to do (or at least would be very unlikely to do) if her information were inaccurate (cf. Goldman, 1999, pp. 108–109). In other words, the author needs to signal that her information is accurate. For example, identifying oneself as the author of a Web site using a *digital certificate* (cf. Froomkin, 1996) can be an effective signal if one

has “such credit and reputation in the eyes of mankind, as to have a great deal to lose in case of . . . being detected in any falsehood” (Hume, 1748/1977, p. 78).

The general properties of such signals have been studied extensively in the literature on game theory (see, e.g., Spence, 1974; Baird et al., 1994, pp. 122–158). For example, just as the quality of information on the Internet varies, the quality of a product typically varies from supplier to supplier. Some suppliers sell high-quality products and other suppliers sell low-quality products. Unfortunately, consumers often cannot tell the high-quality products from the low-quality products just by looking at them. Such situations are referred to as “games of asymmetric information,” where the consumers are the uninformed players and the suppliers are the informed players.

In such situations, suppliers that sell high-quality products would like to send a signal that would allow consumers to distinguish them from the suppliers that sell low-quality products. For something to be an effective signal, however, two conditions must be satisfied: first, for the low-quality suppliers, the cost of sending the signal must outweigh the benefit; second, for the high-quality suppliers, the benefit of sending the signal must outweigh the cost. Under most circumstances, this means that it must cost low-quality suppliers more to send the signal than it costs high-quality suppliers. If the cost is the same for everybody to send the signal, then all suppliers would make the same decision about whether to send the signal.

This game theoretic analysis can be directly applied to indicators of accuracy on the Internet. This is because people seeking information on the Internet are clearly involved in a game of asymmetric information. They often cannot tell just by looking whether the information on a Web site is accurate. Of course, people are not always literally buying the information that they find on the Internet. However, they are deciding how much of their confidence to invest in this information. In addition, the authors of Web sites typically benefit when people invest confidence in their claims.

Using this game theoretic analysis, we can now say more precisely why titular colonicity is not a good indicator of accuracy. Since putting a colon in the title of a paper or a Web site costs everybody the same amount, titular colonicity does not allow people to distinguish sources with accurate information from sources with inaccurate information. Similarly, since automatic spelling and grammar checkers are readily available, the cost of eliminating spelling and grammatical errors is unlikely to outweigh the benefit to an author of having more people believe that her information is accurate.

Notably, changes in technology can have an effect on what will be an effective signal. For example, affordable color photocopying has made it more cost-effective for counterfeiters to include many of the old signals of authentic currency (e.g., green ink). As a result, the United States Treasury has had to come up with some new signals (see Bureau of Engraving and Printing, 2002). Similarly, some of the indicators of accuracy that worked

in the world of print information (e.g., the professional look of a document) may not work as well in the world of electronic information.

VERIFIABILITY OF INFORMATION

Several techniques for verifying the accuracy of information have been discussed in the preceding sections. However, notably, while it will be feasible to verify the accuracy of some pieces of information, it will be extremely difficult (if not impossible) to verify the accuracy of other pieces of information. In other words, some information is verifiable and some information is not.

The distinction between verifiable information and nonverifiable information is commonly made in the literature on game theory. For example, Baird et al. (1994, p. 89) note that "some information is verifiable; that is, it can be readily checked once it is revealed. For example, the combination to a safe is verifiable information. The combination either opens the safe or it does not. Other information is nonverifiable. An employer wants to know whether a guard who was hired was vigilant. The employer might be able to draw inferences from some events, such as whether a thief was successful or was caught, but such information may not be available and, even if it is available, may not be reliable. Even a lazy guard may catch a thief, and a thief may outwit even the most vigilant guard." Basically, a piece of information is verifiable if and only if it is easy to determine whether the information is accurate or inaccurate.²⁹

This characterization provided by Baird et al. captures some of the essential aspects of verifiable information. For instance, the verifiability of a piece of information is independent of the accuracy of the information. The claim that 8-11-64 is the combination to the safe is verifiable whether or not 8-11-64 really is the combination. On the one hand, if you try 8-11-64 and the safe opens, then you have verified that it is the combination. On the other hand, if you try 8-11-64 a few times and the safe does not open, then you have pretty much verified that it is not the combination.³⁰ Similarly, the claim that the guard is vigilant is nonverifiable whether or not the guard really is vigilant. However, a few important subtleties are hidden by (though not inconsistent with) this characterization.

First, this characterization suggests that either a piece of information is verifiable or it is not. However, verifiability is something that comes in degrees. In fact, it is easy to find pieces of information that fall at various different points on the continuum of verifiability. For example, the claim that the safe is gunmetal gray is somewhat easier to verify than is the combination of the safe. Also, the claim that the guard is a descendent of Charlemagne is much harder to verify than is the vigilance of the guard. As a result, we should really speak of the verifiability of information rather than just verifiable versus nonverifiable information.

In general, the verifiability of a piece of information can be measured

in terms of how much it costs to determine whether or not the information is accurate. So, for example, the more time and energy that a person must expend to verify a piece of information, the less verifiable that information is. Of course, in some cases, it may not even be possible to verify a particular piece of information no matter how much time and energy an individual is willing to expend.

Second, this characterization suggests that the verifiability of a piece of information is the same regardless of who is trying to verify it. However, the verifiability of a piece of information depends on the circumstances and capabilities of the individual who is trying to verify it. It will cost you a lot more time and energy to check the combination if the safe is in New York and you are in Los Angeles than it would cost you if you were in the same room with the safe.³¹ Similarly, it is much easier for a trained physicist to verify a claim about the behavior of subatomic particles than it is for a layperson to verify the same claim.

Finally, this characterization suggests that the verifiability of a piece of information is static. However, it is sometimes possible for people to increase the verifiability of a piece of information (cf. Goldman, 1999, pp. 108–109). For example, it becomes much easier to determine whether or not the guard is vigilant if we install a video camera to monitor the guard. As I explain below, this fact turns out to be important as we try to deal with the problem of inaccurate information.

INCREASING THE VERIFIABILITY OF INFORMATION

Library and information scientists have primarily responded to the problem of inaccurate information on the Internet by teaching people how to evaluate information (see, e.g., Alexander & Tate, 1999; Cooke, 1999). Since it is ultimately up to people themselves to decide whether to believe what they read on the Internet (cf. Cerf, 2002), people should certainly receive instruction about what features of Web sites are indicative of accuracy (and about how indicative these features are). Unfortunately, this is not a complete solution to the problem. This is because people do not always apply the techniques for verifying the accuracy of information even when they know how.

People have to act on numerous beliefs to get through their daily lives, but they only have a limited amount of time and energy to expend verifying the accuracy of these beliefs. John Hardwig (1985, p. 335) is not unique when he says that “though I can readily imagine what I would have to do to obtain evidence that would support any one of my beliefs, I cannot imagine being able to do this for all of my beliefs. I believe too much; there is too much relevant evidence (much of it available only after extensive, specialized training); intellect is too small and life too short.” Thus, it is not surprising that Gunther Eysenbach and Christian Köhler (2002, p. 576) found that people seeking health information on the Internet rarely inves-

tigate who the source of the information is even though they believe this to be an important indicator of accuracy.

Fortunately, teaching people how to evaluate information is not the only way that information professionals can try to deal with the problem of inaccurate information. In addition, we can try to make it easier for people to verify the accuracy of information on the Internet.³² In other words, instead of just trying to change the people who are seeking information (by teaching them how to evaluate information), we also can try to increase the verifiability of the information that they seek. If it takes less time and energy to verify the accuracy of information, then people will be more likely to do so (cf. Fallis, 2000, pp. 311–312).

The United States Treasury has adopted just this type of strategy in its efforts to make sure that people can distinguish authentic currency from counterfeit currency. The treasury does not only produce brochures and videos to teach people how to distinguish authentic currency from counterfeit currency. The government also tries to print currency that is easy to authenticate (see Bureau of Engraving and Printing, 2002). Security features (watermarks, color-shifting ink, concentric fine-line printing, polymer security threads, etc.) that are fairly easy for people to detect (and that are fairly difficult for counterfeiters to reproduce) are added to currency itself.

Notably, increasing the verifiability of information may actually be somewhat more cost-effective than teaching people how to evaluate information.³³ Each person who needs to verify the accuracy of a piece of information has to expend energy to acquire and to apply these new skills. However, only one person (e.g., the author) has to expend energy to make the information more verifiable. For example, only one video camera must be installed even though there may be several people who are interested in whether or not the guard is vigilant. As a result, even if the author has to expend more energy to make the information more verifiable than an individual has to expend to acquire and to apply new evaluation skills, less energy will often be expended overall.

TECHNIQUES FOR INCREASING THE VERIFIABILITY OF INFORMATION

Several techniques for verifying the accuracy of information have been discussed in the preceding sections. Information professionals can make information more verifiable simply by making it easier for people to apply these techniques. Of course, it often will be easier for the authors of Web sites to increase the verifiability of information because, like the United States Treasury when it comes to currency, they have more direct control over the information. Also, since the verifiability of a piece of information can depend on the circumstances and capabilities of the individual who is trying to verify it, techniques for increasing the verifiability of information often may need to be tailored to particular audiences. Even so, informa-

tion professionals can do much to make the information that they supply to users more verifiable.³⁴

The basic idea behind making information more verifiable is to give people easier access to evidence (cf. Fallis, 1999). This can be direct evidence of the accuracy of the information, such as the results of empirical studies on the effectiveness of a medical treatment (cf. section 6 above). However, this also can be indirect evidence of the accuracy of the information, such as the medical credentials of the individual promoting the treatment (cf. sections 4, 5, and 7 above).

The main way that information professionals can make information more verifiable is by organizing information (cf. Goldman, 1999, p. 163). For example, this makes it easier for people to find further evidence relevant to the topic that they are interested in (cf. section 6 above).³⁵ Also, this makes it easier for people to consult other sources on this topic to find out if they corroborate the original source (cf. section 5 above).

In addition to organizing information (e.g., by maintaining metadata about the content of information), information professionals can make information more verifiable by maintaining metadata about the context in which information was created and disseminated (cf. Ketelaar, 1997). This makes it easier for people to identify and judge the reliability of the source of the information (cf. section 4 above). It also makes it easier for people to identify and judge the reliability of the methodology that was used to create the information (cf. Giedymin, 1963, p. 289). A related strategy is to promote interactivity on the Internet (cf. Goldman, 1999, pp. 165–166). For example, if the e-mail address of the author of a Web site is provided, people can request specific information that might clear up worries about the accuracy of the information on the Web site.

Another important way to increase the verifiability of information on the Internet is to simply point people toward the accurate information. For example, information professionals have developed portals (or gateways), such as the *Internet Scout Project* (2003a) and *Blue Web'n* (see SBC Pacific Bell 2002), that provide links to Web sites that have been reviewed for quality.³⁶ The fact that these Web sites have been selected by experts serves as an indication of accuracy. As Alison Cooke (1999, p. 39) puts it, “this relieves the user of much of the work in filtering potentially useful sources from the vast quantities of dross available via the Internet” (cf. Burbules, 2001, p. 447). In other words, such portals provide the same epistemic benefit on the Internet that collection management has long provided in libraries (cf. Atkinson, 1996).

Of course, simply creating such portals is not all that needs to be done to increase the verifiability of information. First, such portals need to be publicized. People cannot take advantage of these portals to find accurate information if they are not aware of them. Second, such portals need to publish their selection criteria (see, e.g., Internet Scout Project, 2003b).

Third, information professionals need to develop systems that allow people to “cross-search” several portals at one time (see Worsfold, 1998, p. 1484). Because the construction of such portals is so labor-intensive, a single portal tends to include a fairly limited number of resources as compared with general-purpose search engines (see Worsfold, 1998, p. 1484).

Information professionals also can make information more verifiable by improving the indicators of accuracy. First, we can make it easier for people to check for these indicators of accuracy. For example, Susan Price and William Hersh (1999) have created a computer program that essentially automates the published guidelines for evaluating information on the Internet. The software checks Web sites for the proposed indicators of accuracy rather than a person having to do this manually. This is especially useful for those indicators, such as the number of other Web sites that link to a Web site, that are not immediately observable features of Web sites.

Second, we can create better indicators of accuracy. In other words, we can increase the degree to which existing indicators are correlated with accuracy. For example, an Internet news service, already considered a reliable source of information, might raise its editorial standards. Increasing the degree to which indicators are correlated with accuracy makes it more cost-effective for people to check Web sites for these indicators (cf. Fallis, 2000, p. 311).

In addition to improving existing indicators, we also can create brand-new indicators of accuracy. For example, the HONcode logo was created by the Health on the Net Foundation to be an indicator of accuracy. Only Web sites that abide by this foundation’s code of conduct for publishing quality medical information are allowed to display this logo. In addition, people can simply click on the logo to confirm that a particular Web site is in compliance with the code of conduct.

Of course, simply creating such indicators of accuracy is not all that needs to be done to increase the verifiability of information. First, such indicators of accuracy need to be publicized. People cannot take advantage of these indicators of accuracy if they are not aware of them (or of how much faith to put in them). Second, as noted in section 7 above, such indicators of accuracy must be difficult to fake. For example, unlike many awards for quality on the Internet, it would be fairly difficult for the author of a Web site with inaccurate health information to display the HONcode logo. The Health on the Net Foundation (2002b) has procedures in place that are designed to prevent Web sites that do not comply with their code of conduct from displaying this logo.³⁷

Finally, a couple of somewhat less effective ways of dealing with the problem of inaccurate information on the Internet should at least be mentioned. First, the authors of Web pages can, and probably should, employ practices that make it more likely that their information is accurate. For instance, the author of a Web page might only make claims that have been

corroborated by at least two independent sources (cf. Connell & Tipple, 1999, p. 363). Since there would then be less inaccurate information on the Internet, such practices would make it somewhat less likely that people will be misled. However, such practices by themselves do not make it easier for people to determine whether information is accurate. In other words, simply increasing the accuracy of information is not the same as increasing the verifiability of information. To increase the verifiability of information, the increased accuracy has to be tied to some observable feature of Web sites, such as the HONcode logo.

Second, the inaccurate information could simply be removed from the Internet. In other words, we could impose complete editorial control on the Internet. If inaccurate information is not there, then people cannot be misled by it. Also, in that case, information on the Internet would be easier to verify because the mere fact that the information is on the Internet would be a good indication that it is accurate. However, it is not clear that this strategy would be either feasible (cf. Bruce, 2000, p. 102) or ethical (cf. Doyle, 2001). But in addition, it is not clear that this strategy would really help people to acquire true beliefs. John Stuart Mill (1859/1978) has famously pointed out that there are epistemic costs to restricting access to information. First, since human beings are fallible, some accurate information also would be removed. Second, having access to many points of view (even if many of them turn out to be inaccurate) is essential if people are to acquire justified beliefs. Thus, we should not focus on getting rid of the inaccurate information but rather on making it easier for people to identify it (cf. Wachbroit, 2000, p. 10).

CONCLUSION

As Floridi (2002a) and Herold (2001) have suggested, a lot of work in philosophy is relevant to issues of concern to LIS. In this paper, I have showed how work in the epistemology of testimony is relevant to the issue of how to verify the accuracy of recorded information. In particular, I have argued that this work provides a useful conceptual framework for our efforts in LIS to deal with the problem of inaccurate information. However, it should certainly be noted that epistemology is not the only area of philosophy that is relevant to this issue.

A number of people (e.g., Hardwig, 1994; Shapin, 1994; Burbules, 2001, pp. 450–453) have pointed out that there is an important ethical dimension to the issue of how to verify the accuracy of recorded information. Whenever we rely on someone's testimony, we are putting our trust in this individual. This individual has a moral responsibility not to betray this trust.

In fact, in at least two respects, the moral character of the authors of Web sites is relevant to the issue of how to verify the accuracy of information. First of all, the authors of Web sites certainly have a moral obligation to testify sincerely.³⁸ However, they also have a moral obligation to testify

only when they are in a position to know the fact that they are testifying to (cf. Hardwig, 1994, pp. 88–89). Failing to live up to this obligation is a form of negligence. For example, it might be the case that the author of a Web site “ought to have known” that a particular treatment for a child with a fever was not safe and effective. The problem for people using the Internet is to identify those authors who are living up to these moral obligations (cf. section 4 above).

Of course, it is not just the moral character of authors that is relevant to the issue of how to verify the accuracy of information. Given their special role in selecting and organizing information sources, information professionals also have moral obligations. Whenever we rely on a piece of recorded information, we are trusting that the information professional who has provided that information has done her job. For example, we trust that she has made sure that we will find further information that is relevant to the accuracy of this piece of information, such as a retraction (cf. Fallis, 2000, p. 312). In addition, information professionals arguably have a moral obligation to encourage ethical behavior on the part of authors of Web sites (cf. Hardwig, 1994, p. 91).³⁹

But while ethical considerations are certainly critical to the epistemology of testimony, they are not the whole story. In other words, it is not possible to reduce the epistemology of testimony to ethics as it seems that Steven Shapin (1994) would like to do.⁴⁰ When an individual testifies falsely, it is not always the result of a moral failing. As Peter Lipton (1998, p. 10) puts it, “we may be quite convinced that our informant is an honest chap, yet still wonder whether to believe what he says.” For example, a Web site might promote an ineffective treatment for a child with a fever because of an honest mistake and not because the author has been insincere or negligent. The epistemology of testimony encompasses all of the considerations that are relevant to the issue of how to verify the accuracy of recorded information.

ACKNOWLEDGMENTS

I would like to thank Ken Herold, Kay Mathiesen, and the graduate students in my course entitled *Verifiable Information?* at the University of Arizona for their comments and suggestions.

NOTES

1. For an extensive discussion of the history and scope of the philosophy of information, see Floridi 2002b.
2. Fallis (2000), Goldman (1999), Thagard (2001), Wachbroit (2000), and the articles in the special issue of *Social Epistemology*, 16(1) on “Social Epistemology and Information Science” are some of the exceptions.
3. Such an appeal to work in epistemology is in line with Shera’s (1970, p. 87) call for interdisciplinarity in LIS research. Work in game theory is also an area that Floridi (2002b, p. 139) explicitly had in mind as being relevant to PI. As Floridi (2002a, p. 41) notes, LIS

- does have a broader scope than epistemology. Even so, many of its main concerns are clearly epistemological (cf. Shera, 1970, pp. 82–110).
4. Paul Thagard (2001, pp. 480–481), however, has pointed out that at least some features of the Internet can promote accuracy (e.g., it is much easier to quickly correct errors on the Internet than in print sources).
 5. Ironically enough, library and information scientists themselves have put a lot of inaccurate information on the Internet for the purposes of research (cf. Herson, 1995) and teaching (cf. Piper, 2000).
 6. Giedymin (1963) applies work in the epistemology of testimony to the problem of verifying the accuracy of historical documents.
 7. Locke (1975/1690, p. 663) and Goldman (1999, p. 123) agree with Hume that putting the right amount of faith in the evidence is the best way to maximize truth possession.
 8. Similarly, if the evidence that the claim is inaccurate is greater, then we should be inclined to think that the claim is inaccurate.
 9. The amount of evidence that we should consider depends on how sure we need to be of the accuracy of the information (cf. Burbules, 2001, p. 446). Philosophers typically say that this ultimately depends on what the cost of being misled by inaccurate information is in a particular case (cf. Smith, C., 2002, p. 71).
 10. See, for example, Goldman (1999, pp. 59–65) for a discussion of the correspondence theory of truth.
 11. In addition, philosophers often require that these true beliefs be justified (cf. Goldman, 1999, pp. 23–24).
 12. As Peter Lipton (1998, p. 6) puts it, “the problem of determining which testimony to accept and which to reject . . . presupposes a distinction between what is believed (or at least asserted) and what is the case.”
 13. Bertram Bruce (2000) applies Kaufmann’s notion of “dialectical reading” to the Internet.
 14. For that matter, if our goal is to acquire true beliefs, we also are interested in the accessibility and comprehensibility of the information (cf. Berland et al., 2001).
 15. Goldman (2001) is explicitly concerned with determining which experts you should trust, but what he says applies equally well to determining which Web sites you should trust.
 16. As Hume (1977 /1748, p. 75) might put it, does the information source “have an interest in what they affirm”? Also, notably, bias can result in unintentionally (cf. Goldman, 2001, pp. 104–105) as well as intentionally (cf. Fowler et al., 2001) inaccurate testimony.
 17. Even with all these techniques, it still may be difficult to determine the reliability of a lot of potentially useful sources of information on the Internet.
 18. If someone intentionally gives inaccurate testimony, her goal is not necessarily to mislead. For example, mapmakers often include small inaccuracies to protect themselves against copyright infringement (cf. Monmonier, 1991, pp. 49–51).
 19. Goldman (1999, p. 123) further subdivides this second question into worries about (a) opportunity and (b) competence. In addition, contextualists have pointed out that we might be misled simply because we have a different (e.g., higher) standard of knowledge than the person who is testifying (cf. Smith, C., 2002).
 20. The published guidelines for evaluating information on the Internet do not tend to emphasize this technique.
 21. As Goldman (2001, p. 101) points out, what is required here is only conditional independence and not full independence. If two information sources are reliable, then their reports will be correlated with each other simply because both their reports will be correlated with the truth.
 22. Gwendolen, for example, relies on such direct evidence of accuracy when she says that “their explanations appear to be quite satisfactory, especially Mr. Worthing’s. That seems to me to have the stamp of truth upon it” (Wilde, 2003/1895).
 23. It is interesting to note that, all other things being equal, it is actually the more surprising claims that are more likely to be disseminated. This occurs even with extremely reliable information sources. For example, as Steven Landsburg (1999) notes, “given two papers that have both survived the vetting process, editors [of academic journals] tend to prefer the more surprising, which means that on average they prefer the one that’s wrong.”
 24. In addition, people need to consider how an information source responds to reasons that have been offered by others against the claim (cf. Goldman, 1999, p. 142).

25. This can happen even in science and mathematics. For example, a mathematician with a particular area of expertise may not be qualified to check the proof of a mathematical claim in another area.
26. Cecily, for example, relies on such indirect evidence of accuracy when she says that "I am more than content with what Mr. Moncrieff said. His voice alone inspires one with absolute credulity" (Wilde, 2003/1895).
27. Fallis (2000) shows how these constraints on indicators of accuracy can be formalized using Goldman's (1999) theory of knowledge acquisition. Also, it is important to note that these constraints apply to the issue of who testifies as well as to the issue of how they testify.
28. For each indicator of accuracy that was identified in their empirical study, Fallis and Frické (2002, p. 77) report how much more likely the indicator is to be on an accurate Web site than on an inaccurate Web site.
29. For a piece of information to be verifiable, it does not have to be the case that we can ever know for sure that it is accurate (or inaccurate). With the possible exception of mathematical facts, there will always be some room for doubt.
30. A piece of inaccurate information that is verifiable is sometimes said to be falsifiable.
31. The verifiability of a piece of information can vary with time as well as location (cf. Goldman, 2001, pp. 106–107).
32. In a similar vein, Goldman (1999, p. 108) asks "what can truthful speakers do to make their reports more credible to hearers than they otherwise might be?"
33. Of course, this is not to say that information professionals should not also be teaching people how to evaluate information.
34. It should be noted that, with the exception of removing inaccurate information and creating portals to accurate information, none of the techniques that I discuss in this section requires information professionals to make judgments about the accuracy of specific pieces of information.
35. Goldman (1999, pp. 169–170) also points out how hypertext can allow people to quickly access further information on their topic.
36. Some of these portals, such as the *Digital Library of Information Science and Technology* (2002) and the *Cleveland Digital Library* (see Cleveland State University, 2002), are known as digital libraries, and they collect information resources as well as provide links to other information resources on the Internet. See Cooke, 1999, pp. 34–37 for a list of other portals.
37. Of course, it would be even better if Web sites that displayed this logo without complying with the code of conduct were subject to financial penalties.
38. Not all would agree, however. For example, Gwendolen asserts that "in matters of grave importance, style, not sincerity is the vital thing" (Wilde, 2003/1895).
39. The selection role of information professionals also raises another familiar ethical issue (see, e.g., Wengert, 2001). Namely, when does selection of materials turn into censorship of materials?
40. For an even more radical view that assigns intrinsic moral value to information itself, see Floridi (2003).

REFERENCES

- Alcoff, L. M. (1999). On judging epistemic credibility: Is social identity relevant? *Philosophic Exchange*, 29, 73–89.
- Alexander, J. E., & Tate, M. A. (1999). *Web wisdom: How to evaluate and create information quality on the Web*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Ambre, J., Guard, R., Perveiler, F. M., Renner, J., & Rippen, H. (1997). Criteria for assessing the quality of health information on the Internet. Retrieved September 29, 2003, from <http://hitweb.mitrettek.org/docs/criteria.html>.
- Atkinson, R. (1996). Library functions, scholarly communication, and the foundation of the digital library: Laying claim to the control zone. *Library Quarterly*, 66(3), 239–265.
- Baird, D. G., Gertner, R. H., & Picker, R. C. (1994). *Game theory and the law*. Cambridge, MA: Harvard University Press.
- Berland, G. K., Elliott, M. N., Morales, L. S., Algazy, J. I., Kravitz, R. L., Broder, M. S., Kanouse, D. E., Muñoz, J. A., Puyol, J.-A., Lara, M., Watkins, K. E., Yang, H., & McGlynn, E. A. (2001). Health information on the Internet: Accessibility, quality, and readability in English and Spanish. *JAMA*, 285(20), 2612–2621.

- Bruce, B. (2000). Credibility of the Web: Why we need dialectical reading. *Journal of Philosophy of Education*, 34(1), 97–109.
- Burbules, N. C. (2001). Paradoxes of the Web: The ethical dimensions of credibility. *Library Trends*, 49(3), 441–453.
- Bureau of Engraving and Printing. (2002). Anti-counterfeiting: Security features. Retrieved September 29, 2003, from <http://www.bep.treas.gov/section.cfm/7/35>.
- Cerf, V. G. (2002). Truth and the Internet. Retrieved September 29, 2003, from <http://info.isoc.org/internet/conduct/truth.shtml>.
- Cleveland State University. (2002). *Cleveland digital library*. Retrieved September 29, 2003, from <http://web.ulib.csuohio.edu/SpecColl/cdl/>.
- Coady, C. A. J. (1992). *Testimony: A philosophical study*. New York: Oxford University Press.
- Connell, T. H., & Tipple, J. E. (1999). Testing the accuracy of information on the World Wide Web using the AltaVista search engine. *Reference & User Services Quarterly*, 38(4), 360–368.
- Cooke, A. (1999). *Neal-Schuman authoritative guide to evaluating information on the Internet*. New York: Neal-Schuman.
- Digital Library of Information Science and Technology. (2002). *Digital library of information science and technology*. Retrieved September 29, 2003, from <http://dlist.sir.arizona.edu>.
- Doyle, T. (2001). A utilitarian case for intellectual freedom in libraries. *Library Quarterly*, 71(1), 44–71.
- Eysenbach, G., & Köhler, C. (2002). How do consumers search for and appraise health information on the World Wide Web? Qualitative study using focus groups, usability tests, and in-depth interviews. *British Medical Journal*, 324, 573–577.
- Fallis, D. (1999). Inaccurate consumer health information on the Internet: Criteria for evaluating potential solutions. *Proceedings of the annual symposium of the American Medical Informatics Association*. Retrieved September 29, 2003, from <http://www.amia.org/pubs/symposia/D005366.PDF>.
- Fallis, D. (2000). Veritistic social epistemology and information science. *Social Epistemology*, 14(4), 305–316.
- Fallis, D., & Frické, M. (2002). Indicators of accuracy of consumer health information on the Internet: A study of indicators relating to information for managing fever in children in the home. *Journal of the American Medical Informatics Association*, 9(1), 73–79.
- Floridi, L. (2002a). On defining library and information science as applied philosophy of information. *Social Epistemology*, 16(1), 37–49.
- Floridi, L. (2002b). What is the philosophy of information? *Metaphilosophy*, 33(1/2), 123–145.
- Floridi, L. (2003). On the intrinsic value of information objects and the infosphere. *Ethics and Information Technology*, 4, 287–304.
- Fogg, B. J., Marshall, J., Laraki, O., Osipovich, A., Varma, C., Fang, N., Paul, J., Rangnekar, A., Shon, J., Swani, P., & Treinen, M. (2001). What makes Web sites credible? A report on a large quantitative study. *Conference on Human Factors in Computing Systems*, 61–68.
- Fowler, B., Franklin, C., & Hyde, R. (2001). Internet securities fraud: Old trick, new medium. *Duke Law and Technology Review* (no. 0006). Retrieved September 29, 2003, from <http://www.law.duke.edu/journals/dltr/ARTICLES/2001dltr0006.html>.
- Frické, M. (1997). Information using likeness measures. *Journal of the American Society for Information Science*, 48(10), 882–892.
- Froomkin, A. M. (1996). The essential role of trusted third parties in electronic commerce. *Oregon Law Review*, 75, 49–115.
- Giedymin, J. (1963). Reliability of informants. *British Journal for the Philosophy of Science*, 13(52), 287–302.
- Goldman, A. I. (1987). Foundations of social epistemics. *Synthese*, 73(1), 109–144.
- Goldman, A. I. (1999). *Knowledge in a social world*. New York: Oxford University Press.
- Goldman, A. I. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 63(1), 85–110.
- Griffiths, K. M., & Christensen, H. (2000). Quality of Web-based information on treatment of depression: Cross-sectional survey. *British Medical Journal*, 321, 1511–1515.
- Hardwig, J. (1985). Epistemic dependence. *Journal of Philosophy*, 82(7), 335–349.
- Hardwig, J. (1994). Toward an ethics of expertise. In D. E. Wueste (Ed.), *Professional ethics and social responsibility* (pp. 83–101). Lanham, MD: Rowman and Littlefield.

- Health on the Net Foundation. (2002a). HON code of conduct for medical and health Web sites. Retrieved September 29, 2003, from <http://www.hon.ch/HONcode/Conduct.html>.
- Health on the Net Foundation. (2002b). Policing the HONcode. Retrieved September 29, 2003, from <http://www.hon.ch/HONcode/policy.html>.
- Hernon, P. (1995). Disinformation and misinformation through the Internet: Findings of an exploratory study. *Government Information Quarterly*, 12(2), 133-139.
- Herold, K. R. (2001). Librarianship and the philosophy of information. *Library Philosophy and Practice*, 3(2). Retrieved September 29, 2003, from <http://www.uidaho.edu/~mbolin/lppv3n2.htm>.
- Hume, D. (1977). *An enquiry concerning human understanding*. Indianapolis: Hackett. (Original work published 1748)
- Impicciatore, P., Pandolfini, C., Casella, N., & Bonati, M. (1997). Reliability of health information for the public on the World Wide Web: Systematic survey of advice on managing fever in children at home. *British Medical Journal*, 314, 1875-1879.
- Internet Scout Project. (2003a). *Internet Scout Project*. Retrieved September 29, 2003, from <http://scout.cs.wisc.edu>.
- Internet Scout Project. (2003b). Internet Scout Project report selection criteria. Retrieved September 29, 2003, from <http://scout.cs.wisc.edu/Reports/selection.php>.
- Kaufmann, W. (1977). *The future of the humanities*. New York: Reader's Digest Press.
- Ketelaar, E. (1997). Can we trust information? *International Information and Library Review*, 29, 333-338.
- Kiley, R. (2002). Some evidence exists that the Internet does harm health. *British Medical Journal*, 324, 238.
- Kunst, H., Groot, D., Latthe, P. M., Latthe, M., & Khan, K. S. (2002). Accuracy of information on apparently credible Web sites: Survey of five common health topics. *British Medical Journal*, 321, 581-582.
- Landsburg, S. E. (1999, September 15). Too true to be good: The real reason you can't believe everything you read. *Slate*. Retrieved September 29, 2003, from <http://slate.msn.com/?id=34856>.
- Lee, K. P., Schotland, M., Bacchetti, P., & Bero, L. A. (2002). Association of journal quality indicators with methodological quality of clinical research articles. *JAMA*, 287(21), 2805-2808.
- Lipton, P. (1998). The epistemology of testimony. *Studies in the History and Philosophy of Science*, 29(1), 1-31.
- Locke, J. (1975). *An essay concerning human understanding*. Oxford: Clarendon Press. (Original work published in 1690)
- McDowell, A. (2002). Trust and information: The role of trust in the social epistemology of information science. *Social Epistemology*, 16(1), 51-63.
- Meola, M. (2000). Review of *Knowledge in a social world* by Alvin I. Goldman. *College and Research Libraries*, 61(2), 173-174.
- Mill, J. S. (1978). *On liberty*. Indianapolis: Hackett. (Original work published in 1859)
- Monmonier, M. (1991). *How to lie with maps*. Chicago: University of Chicago Press.
- National Library of Medicine. (2002). Fever and children. Retrieved September 29, 2003, from <http://www.nlm.nih.gov/medlineplus/ency/article/001980.htm>.
- Perry, J. A. (1985). The Dillion hypothesis of titular colonicity: An empirical test from the ecological sciences. *Journal of the American Society for Information Science*, 36(4), 251-258.
- Perry, J. A. (1986). Titular colonicity: Further developments [Letter to the editor]. *Journal of the American Society for Information Science*, 37(3), 177-178.
- Piper, P. S. (2000). Better read that again: Web hoaxes and misinformation. *Searcher*, 8(8). Retrieved September 29, 2003, from <http://www.infotoday.com/searcher/sep00/piper.htm>.
- Plato. (2002). *Charmides*. Retrieved September 29, 2003, from <http://www.perseus.tufts.edu/cgi-bin/ptext?lookup=Plat.+Charm.+153a>. (Original work published in 380 B.C.)
- Price, S. L., & Hersh, W. R. (1999). Filtering Web pages for quality indicators: An empirical approach to finding high-quality consumer health information on the World Wide Web. *Proceedings of the Annual Symposium of the American Medical Informatics Association*, 911-915.
- Rieh, S. Y. (2002). Judgment of information quality and cognitive authority in the Web. *Journal of the American Society for Information Science and Technology*, 53(2), 145-161.

- Salmon, M. H. (1995). *Introduction to logic and critical thinking* (3rd Ed.). Fort Worth: Harcourt Brace.
- SBC Pacific Bell. (2002). *Blue Web'n*. Retrieved September 29, 2003, from <http://www.kn.pacbell.com/wired/bluwebn/>.
- Shapin, S. (1994). *A social history of truth: Civility and science in seventeenth-century England*. Chicago: University of Chicago Press.
- Shera, J. (1970). *Sociological foundations of librarianship*. New York: Asia Publishing.
- Silberg, W. M., Lundberg, G. D., & Musacchio, R. A. (1997). Assessing, controlling, and assuring the quality of medical information on the Internet. *JAMA*, 277(15), 1244–1245.
- Smith, A. (2002). Evaluation of information sources. *The World-Wide Web Virtual Library*. Retrieved September 29, 2003, from <http://www.vuw.ac.nz/~agsmith/evaln/evaln.htm>.
- Smith, C. (2002). Social epistemology, contextualism and the division of labour. *Social Epistemology*, 16(1), 65–81.
- Spence, M. (1974). *Market signaling: Informational transfer in hiring and related screening processes*. Cambridge, MA: Harvard University Press.
- Thagard, P. (2001). Internet epistemology: Contributions of new information technologies to scientific research. In K. Crowley, C. D. Schunn, & T. Okada (Eds.), *Designing for science: Implications from everyday, classroom, and professional settings* (pp. 465–485). Mahwah, NJ: Erlbaum.
- Wachbroit, R. (2000). Reliance and reliability: The problem of information on the Internet. *Report from the Institute for Philosophy & Public Policy*, 20(4), 9–15. Retrieved September 29, 2003, from <http://www.puaf.umd.edu/IPPP/reports/vol20fall00/Fall2000.pdf>.
- Wathen, C. N., & Burkell, J. (2002). Believe it or not: Factors influencing credibility on the Web. *Journal of the American Society for Information Science and Technology*, 53(2), 134–144.
- Wengert, R. (2001). Some ethical aspects of being an information professional. *Library Trends*, 49(3), 486–509.
- Wilde, O. (2003). *The importance of being earnest*. Retrieved September 29, 2003, from <http://www.classicbookshelf.com/library/oscar wilde/the importance of being earnest/>. (Original work published in 1895)
- Wilkinson, G. L., Bennett, L. T., & Oliver, K. M. (1997). Evaluation criteria and indicators of quality for Internet resources. *Educational Technology*, 37(3), 52–59.
- Wilson, P. (1983). *Second-hand knowledge: An inquiry into cognitive authority*. Westport, CT: Greenwood Press.
- Worsfold, E. (1998). Subject gateways: Fulfilling the DESIRE for knowledge. *Computer Networks and ISDN Systems*, 30(16–18), 1479–1489.