
Paradoxes of the Web: The Ethical Dimensions of Credibility

NICHOLAS C. BURBULES

ABSTRACT

THIS ESSAY REVIEWS THE ISSUES SURROUNDING determinations of the credibility of online materials. The author argues, first, that the World Wide Web, and the larger Internet, comprise some very difficult and distinctive features that make conventional ways of assessing credibility adequate only within a fairly bounded frame; second, that beyond this bounded frame, standard credibility measures encounter some paradoxical and self-undermining consequences; third, that this picture is complicated further by the fact that “credibility” actually covers several very different sorts of factors, not all of them matters of judging truth and falsity per se; and therefore, fourth, that the assessment of credibility needs to address the social and normative factors that actually shape the character and quality of online information. These considerations combine to reveal an *ethical* dimension to many credibility assessments.

INTRODUCTION

One of the most-discussed topics about the World Wide Web is how users can be expected to assess the credibility of information they find there. This is not surprising since a key feature of new networked information and communication systems is that the sources of information found online are sometimes difficult to ascertain. The Web seems to offer a global reference resource but, because of its very scope, it seems to overwhelm the ordinary conventions by which people informally judge the merit of what they read or hear. Teaching users how to become more

critical and discerning is an important educational objective for learners of all ages (Bruce, 2000; Burbules & Callister, 2000).

Yet this goal is complicated by the fact that the Web is not an ordinary reference system; it poses some unique and, in many respects, unprecedented conditions that complicate the task of sorting out dependable from undependable information—and even complicates the notion that we have a clear sense of that distinction. How to differentiate credible from fraudulent information is not a new problem, but unraveling these in the context of a vast rapidly changing networked system is.

At a first level, the problems do not seem very different from more familiar text-based or oral contexts. Certainly we are making credibility judgments all the time: Is this person a reliable expert? Does this source have a bias or an axe to grind? Is this information outdated? Does this new information fit with what I already know about a topic? and so on. There are dozens of Web sites already devoted to assessing credibility, and they offer good sensible advice such as: Use the return address or URL to determine the source of the information. Check the “last updated” date to see if the information is current. Triangulate multiple sources of information before you believe something based on what just one source has told you. These are all well and good and, in a large number of cases, will suffice to sort out incorrect, misleading, incomplete, or deceptive information. For learners of a certain age, they are useful rules of thumb, and they are certainly better than nothing. But such standards fail as we consider issues of greater complexity and difficulty, and indeed at some point we realize that they lead us into a series of paradoxes that begin to shatter the notion of “credibility” itself. At that point we are thrown back to much more uncertain tentative methods by which to judge what we find on the Web. Yet this instability itself has something important to teach us about the nature of this new information and communication environment.

Three conditions make the Web, and the larger Internet in which it operates, a different and challenging credibility context. First, there is the problem of sheer volume. A Web search could pull up thousands, or even millions, of “hits” to which one might further add newsgroups, listservs, and e-mail as sources of information on a topic. The numbers are overwhelming. Now, of course, a library can be overwhelming too, as can dozens of news media sources (I write this in the midst of a close presidential campaign and, despite the importance of this subject, it is impossible to find clear unambiguous information on the status of the candidates—each poll gives conflicting numbers, every analysis argues that one or the other has an “edge” in the final election, every assessment of their proposals gives a different calculation of their fiscal costs and benefits, and so on). None of this seems very new. What is new is that the growth and decentered nature of the Web, and the larger Internet, has put the means of providing information in the hands of many more people.

Referencing and organizational systems that are available, for example, in libraries, do not exist here. The markers of institutional credibility and authority, the lines of tradition that allow viewers to judge media sources or publishers, for example, have not been settled yet. There is an even greater capacity to locate information that will tend to confirm one's existing views and prejudices rather than challenge them. In all this, the scope of the network and its deregulated content overwhelm the ordinary idea that we can comparatively judge different sources (which ones?), or that we can trust popular processes of selection to weed out the less credible and give status to the survivors.

Second, the Internet is, to a considerable degree, a self-sustaining reference system—i.e., when we do try to judge the credibility of an information source, we frequently must rely on other information gleaned within the network. We find an article written by an academic group and we go to their university Web site to find out more about them; we find a claim on one Web page and we do a keyword search to see if similar information can be found elsewhere; we receive a rumor via e-mail and forward it to friends and associates to ask if they know whether it has substance or not. Again, often this will help us out and, in any event, it is frequently our only choice. But this self-supporting structure is rather like the problem of dictionaries—i.e., we look up the meaning of one word, it gives us another; we look up the meaning of that word, it gives us another; we look up the meaning of *that* word, and it gives us the first word we started with. When a referencing system operates only internally and has no separate external reference, the very assistance we seek merely leads us in circles within the network. Sometimes credibility judgments online can be like this.

The Web is also a self-supporting reference system in another sense. Because the central feature of the Web is the HTML link, the structure of links by which we access a Web resource, and the links it contains, provides a major source for credibility judgments (Burbules, 1997). On the one hand, how we link to a page usually provides a primary criterion of whether we believe it or not—e.g., it was referred to us by a credible friend or colleague; it was linked to from an authoritative site; several other pages all point to this one reference; and so on. To compare this structure with footnotes (an imperfect comparison generally, but appropriate here): if we find a book or an article cited favorably by other sources we respect, it is more likely to be found important and credible. On the other hand, Web pages typically contain links themselves, and we often judge the credibility of a source by how reliable and complete its references seem to be to other sites. A site for news information that only includes links to other sites expressing a similar political point of view might be viewed as more credible (if one shares that political point of view) or less so (because it only presents one side of the issues). Here too the footnote comparison is fitting—if we know that there are important reference works in a field of

study, we judge new work by how well it anchors its claims to those standard reference points. Credibility here depends on the pathways through which we have accessed the information and the pathways to which it points.

The sort of structure described here is manifested, for example, in search engines that prioritize searches by ranking sites first that are most frequently linked to by other sites. It is manifested in commercial sites that tell, for example, book buyers that: "People who bought this book also bought . . ." It is the core idea behind applications like VisIT, developed at the University of Illinois, which allows one to organize multiple Web sites via their interlinking network structure, ascertaining which seem to be more "central" to an information cluster and which more peripheral (<http://visit1.vp.uiuc.edu/>). This self-supporting nature creates valid opportunities to make some credibility judgments, but it also has the potential to lead those judgments astray, since the closed nature of this reference system may simply mean that we are buttressing one flimsy source by linking it to another one. Unless we have some independent basis for assessing the quality of those associations (whether through our preexisting knowledge of a subject area or through some external-to-the-system standard for judging them), we may be in the position of a blindfolded person being led by the hand by a group of others, all of whom are also blindfolded.

The third factor that makes the Web and the Internet so complicated and difficult from a credibility standpoint is the speed of its growth and the rate of dispersion with which information can circulate within it. Everyone who has spent time with this medium is familiar with the variety of hoaxes, rumors, urban legends, chain letters, and false virus reports that circulate rapidly, often forwarded by users (especially novice users) to their associates, thinking they are providing a useful service by spreading the information. As my colleague Chip Bruce, professor at the University of Illinois at Urbana-Champaign, states it, these become a kind of virus themselves, reproduced geometrically and with great speed. In general, the more important something appears to be, the faster it spreads. Yet when these turn out to be false, the mechanisms for retracting or correcting their misrepresentations can never proceed as quickly or as broadly as the original dispersion. In some cases, this capacity is being used for outright fraud, as in false stories about stock offerings that can often result in substantial short-term increases or decreases in their value; within a short time, even a few hours, such rumors might be detected and corrected, but the savvy crooks have already taken their profits and vanished.

And this introduces one of many paradoxes that impinge upon determinations of credibility online; often the most sophisticated and knowledgeable users are most likely to be taken in by such reports. Believing that they are getting "inside" information not available to the general public, such users are especially prone to being deceived by "obvious"

markers of credibility that—precisely because they are generally reliable indicators of credibility—are easily falsified by clever information-providers with an eye toward deception. A similar example involves falsified Web sites for political candidates, often complete with plausible URLs and apparent testimonials from reliable sources; the truly dangerous ones are not those with obviously satirical or impugning content—e.g., “How I Killed My Brother in Law,” “My Ten Steps for Conquering Canada.” A person with an eye toward real political damage would make the Web site scrupulously accurate, with actual speech transcripts, and so on, adding only a few subtle word changes, a doctored photo or two, a link to a bizarre fringe group, and so on. Disinformation like this actually uses your sophistication against you; the best judges of credibility are most likely to be fooled.

RECOGNIZING CREDIBLE SOURCES

How are we supposed to respond to this situation, and how can we help teach others to be more resistant to it? There are four elements typically discussed in the credibility literature; each has a surface plausibility and range of usefulness and yet each is ultimately inadequate, even self-defeating, as a criterion or procedure.

When examining the many Web sites dedicated to encouraging more critical or discerning consumption of Web information, they have almost identical content (and in many cases have clearly drawn from one another—an interesting credibility problem in itself). The first thing in judging the credibility of a Web site is to look for markers or “proxies” of credibility. These include observing the layout and visual quality of the site (if it is well-designed and carefully maintained, it is more likely to be from a serious source); reading the URL or return e-mail for content (is it from an academic or a commercial source, does the provider identify him or herself by name); does the information seem to come from an authoritative source or one with an obvious bias; does the date on the material show that it is “fresh” and frequently revised and updated; and so on. These indirect indicators have value but, as we have already seen, deceivers are as aware of them as consumers are, and each can be falsified not only with an eye toward casual deception but as a way of taking in people even more profoundly. Moreover, these criteria, though widely shared, are far from unambiguous. Is an academic (.edu, .ed, or .ac) URL marker a clear sign of credibility? Academics tend to think so, but many people view it differently, seeing academic status as a marker of irrelevance, abstraction, or arcania. There are also issues on which an academic identification might be viewed as a sign of special pleading by a self-interested constituency (e.g., tax policies) or of Left-wing partisanship. In other words, these criteria themselves rest upon additional judgments that are often tacit or unexamined, without which they may be extremely unreliable guides to judging information online. Yet the further one examines these

unquestioned assumptions, the more flimsy these criteria appear. Given the value that academic institutions place on the originality of ideas, there is a strong incentive for scholars to look for the unexamined angle of approach to a problem, the radical interpretation, the obscure detail elevated to a masterstroke of brilliant insight, the replicated study that disproves a long-established “fact,” and so on. Certainly academic practices impose rigorous standards of scholarship, peer review, critical cross-testing by other scholars, and so on. Academic journals are one of the media that promote and protect such values. But there are simultaneous pressures toward novelty and academic status or visibility that also operate for many scholars, and which do not always promote participation in the critical spirit of inquiry that presumably guarantees the credibility of academic work. None of this is apparent to those outside the academic context, so when an article comes from “Dr. Smith from Recondite State University,” its worth rests on an invisible set of norms and practices that even most academics tend to trust on faith; ordinary users don’t even know that those practices exist, let alone how imperfect they can be.

A second set of responses, then, takes the opposite tack: be skeptical of everything found online; use multiple indicators of credibility, including those *external* to the source and not only those identified within it; triangulate specific claims by matching them with information available from independent sources (whether online or not); in general, do not believe anything that comes to you from only one source. Again, these are perfectly reliable rules of thumb. But they probably tend to exclude too much, and they require a degree of thoroughness that, realistically, few people will apply to every case. Determinations of credibility are not a perfectionist endeavor; they inevitably entail judgments about *how much* credibility one needs to support action or belief based on a particular claim, and this degree will generally vary depending on the seriousness of the consequences for that person of an error. But there are two kinds of error at work here (statistically termed Type I and Type II errors): the consequences of rejecting a true proposition can be just as devastating as the consequences of accepting a mistaken one—and nothing in lists of credibility criteria can help in balancing *that* determination. More stringent criteria may decrease the likelihood of making one kind of mistake at the cost of increasing the likelihood of making the other kind of mistake.

A third response, and probably one that most of us follow most often, knowingly or not, is simply to defer these judgments to others whom we entrust to make them on our behalf. Even search engines are making qualitative judgments about sources (different engines use different criteria), which are implicit in the rank ordering they establish when they post the results. Many users may not know that some search engines “sell” priority in their listings so that the criterion is based on commercial con-

cerns and not necessarily on the quality, reliability, or usefulness of the sites listed. Directories, whether partly automated or driven by human editors/archivists, often seek to establish “definitive” sites, selecting, evaluating, and organizing exemplary resources on some topic, and granting through this process a kind of derivative credibility (assuming, of course, that these editors/archivists are reliable judges of material themselves). I have already touched on the crucial role played by editors and publishers of journals, or other kinds of online publications, in screening and authorizing particular works as worthy of attention. All of these functions are perfectly recognizable to librarians, of course. It is far from an original insight to note that this whole process simply raises the question of credibility at a level once-removed. Yet it is important to emphasize here that it is often difficult, if not impossible, to identify who these intermediaries actually are: it is natural to want the primary material to stand for itself and, by disposition, people in these sorts of roles are not interested in interjecting themselves as the focal points of attention. But as I hope to have made clear, without doing so, users are deprived of information about a crucial element in the credibility chain. At the same time, and in a manner similar to my second case mentioned earlier, the more “meta” such reflections become, the more that judgments on the credibility of the intermediaries and facilitators of access to information displace judgments on the credibility of that information itself, the more time users spend away from the things that they actually want and need; many will decide that it is not worth the time or will use very broad imperfect standards (e.g., “refereed publication”) to perform a kind of “information triage” for them—rapidly, crudely, but necessary given the volume of material to be worked through.

A fourth related approach—one distinctly suited to the Web—is the formation of communities (“rings”) of like-minded people who share a common interest or concern. By linking their Web sites together and collectively screening the addition of new material, they pool their intelligence and expertise to make credibility judgments and to cross-check one another. This phenomenon is interesting both as an epistemic exercise and as an instantiation of social constructionism at work. However, obviously it is imperfect since shared wisdom can also mean shared misconceptions or biases. While less hierarchical and more democratic than relying on invisible editor/archivists to make judgments on one’s behalf, this approach has the vices of its virtues. One might term this an instance of “distributed credibility” in that it displaces an individual judgment with a collective intelligence. It is, as I mentioned, particularly “Web-like” in its holistic approach to knowledge but also in its self-supporting and potentially self-reinforcing character. The greatest danger of such communities, as with communities generally, is that they can become exclusionary, hostile to unconventional, or radical challenges to their presumptions and

practices (Burbules, 2000). From a credibility standpoint, this means that serious questioning—the kind of questioning that can only come from one “outside” a given epistemic framework—is less likely to occur, and it is more likely that over time the shared preconceptions of such communities, even when they may have been originally valid, will eventually become credibility blinders.

What I have tried to show here is how the most common responses to credibility issues online, while valuable and reasonable within certain constraints, ultimately turn out to be paradoxical and self-defeating. This does not make them useless, but it suggests a limit to how clear and reliable such credibility judgments can be. At some juncture they encounter a point of diminishing returns or, as Tenner (1996) calls them, “revenge effects” that actually counteract one’s purposes (pp. 5-6).

DEFINING CREDIBILITY

This discussion is complicated still further by the fact that “credibility” means many different things, not only in the Web context but also generally. Normally it is taken as synonymous with “truth” or “believability” and is tied together with the epistemic problem of how *information* becomes *knowledge*. Certainly, the most striking examples of rumors and hoaxes online concern actual misinformation or disinformation intended to deceive others. But these concerns, important as they are, should not obscure other key dimensions that also impinge on judgments of “credibility.”

Such judgments also depend on assessments of what is *useful*, *relevant*, or *interesting* (and these are not all the same thing either). Given the volume of online information, a major decision to be made is simply what to pay attention to and on what basis. The initial selection and screening, which I would call “the judgment of credibility,” is typically based more on one’s interests and concerns and whether this new item even potentially qualifies as worthy of attention. Giving over one’s attention is what starts the process of epistemic evaluation, but this way of putting it suggests that these are entirely separate processes. In practice, we may have already made tacit preliminary judgments of truthfulness as soon as we say, “this may be useful, relevant, or interesting.” Conversely, it may partly be because we find something potentially useful, relevant, or interesting that we have an incentive to find it true.

When we look at the role of intermediaries in the credibility process, we need to scrutinize them not only for their qualifications and their criteria for making epistemic judgments on our behalf but also, crucially, *their* judgments about what is useful, relevant, or interesting. This latter role is often underestimated, as can be seen, for example, with the news media (whether online scribes like the Drudge Report, or CNN and the BBC). The decisions they make about what and how much to tell us about

certain stories are based only partly on the value of truth—e.g., operating with limited time and a need to grab and hold viewers' attention, sharp cuts are made in stories with an eye toward what viewers will find useful, relevant, or interesting. The *New York Times* slogan "All the News That's Fit to Print" becomes "All the News That Fits," often in a few-second sound bite. So here again the different dimensions of credibility cannot be divorced from one another. This becomes even more true when we reflect on the feedback effects of such decisions by the media (whether online or not)—i.e., their judgments about what people will find useful, relevant, or interesting often shape, at least in part, what people actually *do* come to find useful, relevant, or interesting. This self-fulfilling influence makes it not only a consumer-driven determination but a consumer-shaping determination. And when these media are seen in the context of their increasingly intrusive commercial interests, this process can hardly fail to be viewed skeptically.

Another dimension of credibility is *timeliness*. This is not only important in the ordinary sense that much information becomes obsolete or inaccurate with the passage of time, it is also because the Web is such a rapidly changing environment that material which does not appear to be continually revised and checked becomes suspect for no other reason than that the environment around it has been changing. Hence, many users expect a degree of ongoing novelty and innovation even just at the level of design or mode of representation, as an indication marker that the providers of information have been scrutinizing their content with a "freshness date" in mind. This is hardly a new idea. In the novel *The Leopard*, Giuseppe di Lampedusa (1961) writes:

a fact has scarcely happened five minutes before its genuine kernel has vanished, been camouflaged, embellished, disfigured, annihilated by imagination and self-interest; shame, fear, generosity, malice, opportunism, charity, all the passions, good as well as evil, fling themselves on the fact and tear it to pieces; very soon it has vanished altogether. (p. 219)

Another dimension of credibility is *comprehensiveness*. The very volume and diversity of the Internet creates a peculiar credibility dilemma. One might term this Meno's Paradox in reverse: How do you know what's not there when you do not know what it is? The global scope of the Web can create the illusion that whatever cannot be found must not be very important. Yet even if everything one did find were important and true, it would be a significant failing from the standpoint of credibility if other information, representing contrasting or conflicting points of view, was not also available. In part this is because, as noted earlier, having contrasting or conflicting points of view is sometimes the only check on a kind of self-confirming "obviousness" that beliefs can settle into, but it is also because many of the intermediaries, as we have seen, play their primary role through

selectiveness, a role we implicitly authorize for them. We *want* others to make decisions about priority and relevance; we *don't want* the "full story" (whatever that means) in most cases. It would be tedious and distracting otherwise. But as soon as such selections are made—by others or by ourselves—the danger increases that something crucial has been overlooked. And without substantial independent knowledge of a subject area, it is impossible to find out, even with hindsight, what has been overlooked. The fact that such selections are absolutely necessary in the Web context only heightens the paradox: the selectivity we require for making certain kinds of credibility judgments conflicts with the comprehensiveness we require as a condition of other credibility judgments.

Finally, these considerations lead to another paradox of credibility, since sometimes *too much* comprehensiveness can itself be counterproductive to judgments of credibility. Credibility, as noted earlier, involves judgments about worth and not just about truth per se. David Shenk (1997) argues that an excess of online information tends to produce a "leveling" effect in which everything is viewed as equally plausible or implausible—i.e., for any point of view you can find a reasonably well-argued alternative view. How does one decide then which perspective to valorize? When reasonable arguments seem to pull in any one of several directions, does it matter which one we choose? This judgment typically will rely on value considerations beyond the force of argument itself; credibility is not a purely epistemological assessment.

CONCLUSION

What this discussion has tried to show is that the standard criteria for judging credibility online are frustrated by the characteristic conditions of the World Wide Web and of the larger Internet. None of these elements is entirely unique to the online context, but the scope, self-referencing character, and rate of change of this medium raise these issues to a new importance. The Web is both an information archive and a social network; as people move within this space, their interaction with ideas and information is, at the same time, an interaction with other individuals or groups (even though the implications of such social dynamics for what is and isn't found online are not always made apparent). The analysis here has tried to make such implications more apparent. Moreover, credibility is not just one thing, and judgments about it inevitably bring in considerations that are not purely matters of assessing knowledge claims. At this point credibility can be seen to take on an *ethical* dimension.

The idea of distributed credibility suggests that the reliability of judgments about the truth of information—and even more so judgments about usefulness, relevance, interest, or worth—cannot be assessed outside the nature of the online communities of which one is (overtly or tacitly) a part, nor of the communities producing and legitimizing the information

found online. These collective sentiments, with all their wisdom and insight, all their biases and exclusions, shape the content of information, shape the standards by which it is judged, and shape the negative spaces, the absences, of what is not to be found there. These social judgments and processes are structured by values of collectivity, reciprocal obligation, inclusiveness or exclusiveness, and so on—moral values, not primarily epistemic ones. Hence judging information is always partly judging other values with which one is choosing to identify or to challenge. A key element in these judgments, I have suggested, is the permeability of such communities to questions or challenges from alternative points of view; some, like gated communities in our physical neighborhoods, define themselves centrally by what they *exclude*; others, like communities of open debate and free inquiry, invite and even seek out alternative perspectives because they believe that this is how knowledge is best formed and tested—but also, and inseparably from this belief, because this is how they choose to shape the world in which they live. That is a moral stance as well.

On the issue of what I described as Type I and Type II errors, no epistemic criteria are going to inform the decision about whether it is better to run the risk of mistakenly accepting a false proposition or mistakenly rejecting a true one. Such a judgment cannot be made once and for all time; the importance of the problem, the consequences of error, and upon whom those consequences may fall, all impinge on the way this choice will be decided in particular cases. Determinations of risk, in such instances, clearly raise moral questions, and it is crucial to see that, whether decided overtly or taken for granted, judgments of this sort implicate the process of inquiry in a process of social responsibility.

On the issue of how “meta” our reflections on online information should be—that is, how deeply to investigate the sources of information, their qualifications, their procedures of testing and confirmation for what they put out as fact—the ideal of perfectionism in such matters runs up against realities of limited time and limited available information. At some point we need to move back down from the meta level and get the information we were originally looking for. This means that, at some point, judgments rely on attributions of trust—trust in individuals, trust in communities and collective (if largely invisible) processes of vetting information. Here again we encounter moral elements: fostering such relations creates, in networked environments especially, fibers of affiliation that strengthen and reinforce one another. For example, if one visits a Web site frequently for information because one judges it to be reliable, one contributes to granting it a higher priority when search engines rank it for the searches of others—they are bound up with one’s own choices, whether they choose to be or not, and their participation with these sites, in turn, reinforces the *reputation* of credibility that the site acquires. This may not have been one’s intention, but the unspoken, and unquestioned,

attributions of trust that are implicit in these networked dynamics introduce a problem that demands moral reflection.

On the issue of "information triage," the notion of epistemic perfectionism takes an even stronger hit: credibility models that are based on the ideal of scrupulous testing and comparison fall up against practices of inquiry that, for most users most of the time, simply do not work that way and cannot. The need for selectivity, and the general expectation of speed and efficiency in online searches, dictate that users either make fairly rapid choices or leave those determinations up to others. The criteria by which such accommodations get made involve all kinds of assumptions, preferences, and blind spots—but any alternative set of criteria would also involve assumptions, preferences, and blind spots, simply different ones. Deciding how much time to spend on credibility testing, then, is not itself a credibility issue but a matter of values and priorities that bring in other considerations (for example, I am typing the first draft of this paper with my infant son asleep on the couch next to me; even as I make revisions and try to improve the paper's arguments, its credibility, it will always bear the marks of having been written with one ear perked, listening for his movement or discomfort—the same would be true if I were searching for information online).

On the issue of timeliness and the erosion of "facts," I recommend an exercise for you, which comes from Bob Panoff of the Shoder Education Foundation: search the Web for the boiling point of radium. You will find dramatically different figures, all from sites that appear legitimate and credible (the first four I found had four different temperatures). Since there is no way (I assume) for you to ascertain the fact of the matter for yourself, how will you decide which site to believe? If this is true for something that is assumed to be (a) an objective and scientific fact and (b) constant, since whatever the boiling point is presumably isn't going to change, how much more difficult will it be for other sorts of "facts" to be verified? For information with a stronger social or political dimension, it is clear that variations in empirical claims will be inevitably wrapped up with social or political values or assumptions, and where these claims conflict, as they inevitably will in a vast, networked, information environment, how could normative values and commitments *not* be a factor in deciding what/who will be believed?

Finally, on the issue of what Shenk (1997) calls the leveling effect of too much and too eclectic a pool of information, the greatest danger, as noted, is that one will simply choose to accept information that plausibly confirms one's prior beliefs or what one wishes were true. None of us can be entirely immune to this weakness but, to the extent that credibility judgments are recognized as having an ethical element, the consequences of doing so, for ourselves and for others, can at least be brought to the surface.

And that, in the end, is what I am arguing for here. The notions that credibility judgments can be made on objective criteria, that they only involve considerations impinging on the truth or falsity of information, and that one should always exercise such judgments as scrupulously as possible, all neglect the underlying characteristic of the networked environment in which these judgments are being made, namely, that it comprises—constitutes, even—communities of obligation and commitment. The social dimensions of this network always entail elements of judgment and value. In the end, the best safeguard is to check one's judgments against the judgments of a community with which one has confidence; choosing that reference group prudently is as much a moral matter, involving issues of respect and trust, as a matter of expertise.

As noted throughout this discussion, these ethical elements, along with the fundamentally complex and conflicted nature of what judgments about "credibility" actually involve, combine to make such judgments far more indeterminate and provisional than they are normally taken to be. In a context where epistemic perfectionism must take its place alongside many other values that frequently override it, credibility, in the sense it is often meant, often may not be able to be independently established at all.

REFERENCES

- Bruce, B. C. (2000). Credibility of the Web: Why we need dialectical reading. *Journal of Philosophy of Education, 34*(1), 97-109.
- Burbules, N. C., & Callister, T. A., Jr. (2000). *Watch IT: The risks and promises of information technologies for education*. Boulder, CO: Westview Press.
- Burbules, N. C. (1997). Rhetorics of the Web: Hyperreading and critical literacy. In I. Snyder (Ed.), *Page to screen: Taking literacy into the electronic era* (pp. 102-122). London: Routledge.
- Burbules, N. C. (2000). Does the Internet constitute a global educational community? In N. C. Burbules & C. Torres (Eds.), *Globalization and education: Critical perspectives* (pp. 323-355). New York: Routledge.
- Lampedusa, G. d. (1981). *The leopard*. London: Fontana Books.
- Shenk, D. (1997). *Data smog: Surviving the information glut*. San Francisco: Harper Edge.
- Tenner, E. (1996). *Why things bite back: Technology and the revenge of unintended consequences*. New York: Knopf.