

Reliability on the Crowded Net

Finding the Truth in a Web of Deceit

P.D. Magnus
pmagnus@fecundity.com

"Information isn't knowledge."
—an ad for Fortune magazine

Abstract

On-line, just as off-line, there are ways of assessing the credibility of information sources. The Internet, although it arguably makes for nothing wholly new in this regard, complicates the ordinary task of assessing credibility.

In the first section, I consider a specific example and argue that Internet content providers have no clear interest in resolving these complications. In the second, I consider four general ways that we might assess credibility and explore how they apply to life on-line. Finally, I argue that even careful information gathering may pose subtle pitfalls but that some of these may be mitigated on-line.

On the Internet, everyone is encouraged to have and express an opinion: On the evening news, an anchor encourages me to visit the station's web site and share my opinion about the day's news. At the same time, the Internet has become a first line of research: If I wonder what the capital of Mongolia is, for instance, I am as likely to search the web as I am to search my shelves for an appropriate book. The desire to find facts on the Internet is potentially in conflict with the proliferation of information: If everyone on the net were to take their best guess as to the capital of Mongolia, I would be ill-advised to trust some random selection from the responses.

The Tension Between Expression and Reliability

To take one example, Amazon.com encourages visitors to post reviews of the products that are available for purchase from the site and even encourages visitors to vote whether reviews written by other visitors were helpful or not. If Amazon were interested in furthering the cause of knowledge, perhaps it might be relied upon to disarm the potential conflict

between inclusiveness and reliability. Since Amazon is interested primarily in selling books (plus toys, stereos, and so on), it can only be expected to promote reliability if that in turn promotes its bottom line. Does it? It could be argued that accurate reviews make for happy customers, happy customers will return to buy more, and thus Amazon does have an interest in discouraging false or inaccurate information. This supposes that the reviews at Amazon's site are primarily a service to customers, and this is not obviously the case. Encouraging visitors to write product reviews encourages them to associate themselves with Amazon's site and thus promotes customer loyalty. There is another, perhaps less obvious, reason.

Consider that Amazon, like any business, has an interest in collecting information about its customers. By using cookies, Amazon can track where on the site a visitor goes, but the information that can be gathered with cookies is incomplete for a host of reasons. Users may visit the site from more than one computer, multiple users

may share single computers, and visitors may refuse to accept cookies.

In order to enter a review, visitors must identify themselves by entering their Amazon user name and password. A visitor must do this even if only to rate a review as helpful or not helpful. There are rationales for this requirement; for instance, it assures that a user does not vote a dozen times for the helpfulness of his own review. Yet it also allows Amazon to associate users with particular product preferences, offering many advantages over cookies alone. Users voluntarily supply information that identifies them uniquely. Where a cookie may allow Amazon to determine that a particular user looked at the product description for a digital camera and a stereo system, if that user rated several reviews for the camera but none for the stereo then Amazon may infer that the user was more seriously interested in the former than the latter. Information gathered from reviews may be combined with cookies, of course, to provide more marketing information than either method alone. Thus, Amazon can gather considerable marketing information about its visitors from the system of reviews, regardless of whether the reviews spread accurate or inaccurate information. The point is not that Amazon will make insidious use of this information or that it represents an invasion of visitors' privacy; perhaps they will and perhaps it does, but perhaps not. The point is that the marketing value of the system of reviews is a sufficient reason for Amazon to encourage it.

Given that allowing reviews costs Amazon very little, that policing them for accuracy would increase the cost considerably, and that it promotes Amazon's interests to have the reviews regardless of their accuracy, one can hardly expect Amazon to assure that the information provided in reviews is accurate. One may easily shrug off this analysis, since it concerns reviews.

Reviews, one might think, are matters of opinion. Whether or not Joe in Tuscaloosa enjoyed the latest Tom Clancy novel is not the sort of fact anyone searches the Internet to discover. Moreover, if the review only addresses issues that are matters of taste, then it is not clear that Joe can be wrong. He may say things with which I would disagree, but this would only be because our tastes differ. On the Internet, just as off-line, there is no disputing matters of taste.

The reply here is threefold. First, reviews are not always *mere* statements of taste. In a review of the movie *The Seven Samurai*, "A viewer from Akron, OH" states:

Lest anyone be dissuaded from purchasing this masterpiece because they believe it is not presented in its original aspect ratio, it should be known that THE 1.33:1 ASPECT RATIO ON THIS DVD IS CORRECT. Akira Kurosawa did not begin working with the widescreen format until later in the 1950s. Anyone who asserts otherwise is mistaken.

The review makes a factual claim, the truth of which is critical to its value as a review. Moreover, it contains encouragement to ignore any competing factual claims.

Second, claims which we read but had not sought out are sometimes the most insidious. Take the example above. Sometime after reading the review, I found myself in a conversation in which I offered the claim as if it were a fact. Because I was not concerned with the aspect ratio of the film at the time I read the review, I did not critically evaluate the source. At a later time, when I was concerned to know the aspect ratio, I only remembered that it was not widescreen. In this case I am able to reconstruct where I had acquired that belief— although I suspect I may have read it somewhere else as well— but it

is probably the exception. Claims which are trivial to us when we read them may easily lay claim to our belief, but we often continue to believe them even when they become of some consequence.

Third and perhaps most importantly, the phenomenon is not confined to reviews. Sites like Slashdot.org, Kuro5hin.org, and Im-Ur.com offer stories meant as items of news, but written or submitted by members. These sites rely on members for submissions and on those submissions being interesting to visitors, but they do not obviously rely on the truth or reliability of the news items they carry. Further, the low-cost of hosting allows any crackpot to post his delusions to the web— and low-cost hosting services encourage him to do just that.

The tension between free expression and reliability is inevitable in a free society, and we have ways of dealing with it off-line. Even if life on-line poses no fundamentally new problems in this regard, it exacerbates problems and disables old coping mechanisms. In the next section, I consider several methods for assessing credibility on-line.

Four Methods of Assessing Credibility

So the quantity of content on the Internet results in a decreased reliability for arbitrary selections from that content. We are thus confronted with the problem of determining which claims should be believed. The most obvious way of determining this is to look for reliable sources. This shifts the problem from deciding which claims are to be believed to deciding which sources are to be trusted.

Appeal to Reliability

We may appeal to our background beliefs to decide whether someone will reliably make accurate claims. As an example, suppose someone— let's call him E— goes to a fan site that describes

what happened at a recent concert. Suppose that the fan site is maintained by H, who claims to have been at the concert herself. If E believes that she was at the concert and that she is being honest, he is justified in believing that things happened as she describes. His ability to trust a particular claim (that such and so happened at the concert) relies on his accepting an indefinite list of other, background claims (that H went to the concert, that H is not spreading malicious rumors about the band, and so on). Call this the method of *appeal to reliability*. The problem with it is readily apparent: how can E be sure of his background beliefs?

Off-line, he has ways to address this question. H may be a friend of E or at least a friend of a friend; someone besides H herself will swear for her. Perhaps H can show her ticket stub to E or provide other physical evidence. Such checks allow him to judge whether she went to the concert and so on; he can make these checks against a source besides her own testimony.

In an on-line community, chat room, or IRC channel, E may be able to rely on similar resources. E may trust some members of the group who would swear for H. This is only possible, though, if E has been involved in the community long enough to trust some members of it and if H has been involved long enough that those others will swear for her. If E only knows H from her web site, it may be that all he knows about her is what she has had to say for herself.

Although illustrated in terms of individual people, these issues arise with institutional sources as well. Consider, for example, that Billy Barty died last December. Suppose E believes this because he read it in the New York Times and because he believes that the Times is typically accurate about this sort of thing. It makes little difference

whether he read it in the actual paper or on the Times web site.

Some news sites— like nytimes.com— are extensions of traditional publications, and we would trust them to whatever degree we would trust their traditional counterpart. In the proliferation of news web sites, however, many do not have an off-line sponsor from which to derive credibility. Evaluating the reliability of such sources can only be guided by their on-line presence, but the very question is whether their on-line claims are reliable. This circularity makes appeals to reliability ultimately unsatisfactory.

Appeal to Plausibility

If we are unable to determine which particular sources are to be trusted, there are still ways that we can cull through content. We might assess the plausibility of a particular claim, independently of whomever the source might have been. In addition to judging content, we might judge the plausibility of the form in which a claim is presented. To illustrate the former: If E reads on a web site that Princess Diana's death was arranged by a conspiracy of Arab oil men, then he may reject it on grounds of plausibility. To illustrate the latter: Suppose E does a search for high-energy physics research and finds the web page of someone, call her P, who claims to have a PhD in physics and who describes several cutting-edge projects. How should E evaluate what P has to say? If P does not write using the language of a professional physicist, E may doubt whether P had really completed a PhD. If P writes in the style of a high school student, E should probably not give the site much credence. Note that E can make these judgements regardless of what particular claims P is making. Call this method, whether applied to content or to form, *appeal to plausibility*. Unfortunately, it only yields negative results. We should be disinclined to believe things that are implausible, but

the fact that something is plausible does not by itself show that we should believe it.

Calibration

E still has other ways to assess the content of P's web page. He can compare the partial knowledge he has of high-energy physics with the claims made on the web page. If the page is correct on those points, then the page might be expected to be correct on other points as well. Call this the method of *calibration*: E can judge the page reliable if, on the things he can check, it gets things right.

The problems with calibration turn on which facts E is able to check. Suppose P's website contains some claims taken from a standard physics textbook and others from P's fevered imagination. If E reads the standard textbook to check the claims made by P, then he will find that every claim he can check is correct. If he supposes on the basis of this that P's site is generally trustworthy, he will accept P's delusional ravings. If E's base of knowledge is not systematically related to P's, though, the method of calibration should give E a good measure of how much to rely on P's web page.

Both appeal to plausibility and calibration require that we have some non-trivial knowledge about the subject matter that interests us before we can evaluate on-line sources on the topic. This difficulty is magnified because we often go on-line to look for information about something *just because* we know little about it ourselves.

Sampling

These methods consider a particular report or web page in isolation, but it is further possible to assess the reliability of a claim by comparing related claims made by several sources. If all sources disagree, then the choice of which if any to trust must be made on the basis of some other method. If all or most agree, then those attract more credibility than

they would if they were considered separately. This method, which we may call *sampling*, will correct for individual bias and error but will fail to increase reliability if distorting biases are shared by all or most of the sample.

A mundane example of sampling is the practice of asking for a second opinion. If E is diagnosed with a serious illness and his doctor recommends surgery, he may insist on asking another doctor. If the diagnosis was the result of some mistake or if his doctor was unusually keen on surgery, then the second doctor will hopefully offer a different opinion. If the test used to detect the disease systematically returns false positives or if both doctors rely on the same protocol to decide if E's problem requires surgery, then the second opinion will accord with the first and lead E to place potentially unwarranted trust in the diagnosis and recommendation.

Sampling has similar strengths and weaknesses when used to evaluate claims found on-line. If a website contains errors that result from authors being careless, confused, or idiosyncratically deluded, other pages are unlikely to contain those same errors; so, a comparison of multiple pages can separate those errors from other claims. Conversely, where many authors have copied information from a common source or from each other, sampling will not increase reliability. Even if authors do not share sources, sampling will fail to help if they share some distorting bias. Sampling applied to claims made on the Internet automatically excludes people who do not use the Internet, and for some topics this may introduce a significant bias.

Combinations

The four strategies discussed here exhaust the ways in which a claim found on the Internet can be reasonably evaluated. They may, of course, be used in conjunction with one another. Sources that calibrate positively may

count for more in sampling; claims made by one reliable source may be used as a basis against which to calibrate another source; and so on. In cases where the methods are insufficient to justify much confidence, we are forced either to look to more traditional sources or to do without knowing what we had gone on-line to learn.

Permutations

The four methods considered above address the question of whether or not a claim found on-line is to be believed, but variants of them may be used to *infer* what to believe from what can be found on-line. Consider a discussion earlier this year on ArsTechnica.com which degenerated into a flame war between people who insisted that AMD chips are great and others who insisted that AMD chips are junk. One might conclude that AMD chips outperform but are less reliable than Intel chips. This conclusion isn't stated explicitly by either party in the heated exchange, but may be inferred from the fact that many of respondents report high performance from the chips along with the fact that many respondents report that the chips can not be made to work. This is something like sampling, but the conclusion of it is not to trust either party of flammers.

Pitfalls of Information Gathering

The four methods make it possible to assess the credibility of a claim found on-line, but problems remain even if we are vigilant; a claim may be misleading even if it is true. If some contributor to a discussion board complains about a particular product, then people who read that post may well be left with a poor impression of the product and its manufacturer. Yet the complaint— even if true— represents the experience of only one of thousands of customers. Ideally, a sample of different posts about that product will give a fair distribution of responses. However, it will often be the case that unsatisfied customers are more likely to write about

their experiences than satisfied customers, while for some products it may be that zealous advocates for the product are most vocal. In either situation, judging the reliability of this or that claim is not enough. It is also essential to determine how the contributors' experiences with the product reflect on the product generally.

In addition to uncertainty about the veracity of claims, then, there can be uncertainty about the significance of them. This may occur with traditional media too, of course. Increasingly prominent coverage of crime in a newspaper or on television— even where factually accurate— can foster the inaccurate perception that crime is on the rise. There is some hope that there is actually less of this sort of misrepresentation on the Internet than in traditional media. Slate.com, which summarizes and condenses news from around the world, indicates where different newspapers print stories. It thus reinforces the lesson that there is disagreement about the lead story on any given day. It is easy enough to download the front pages of several major newspapers and make this sort of survey yourself. Moreover, the low cost of putting together a webpage allows views to be represented which would not be represented prominently— and perhaps not at all— in traditional media. This is just the feature of the Internet that causes problems for reliability. This should come as no surprise: For any claim, free expression increases the likelihood that someone will make it. This holds for true claims as well as false ones.

Conclusion

Cybernetics began as the study of information systems, so it should be no surprise that cyberspace opens new possibilities for gathering information and presents new challenges to gathering reliable information. There are ways to meet the challenge, and they are the same ways we have for gathering

reliable information out in meatspace. In a sense, then, this paper has merely rehearsed things we already knew. There is some value in making them explicit, however, and more in holding them up to reflection.