Chapter 9

Trolls, Spin, and the Boundaries of Trust

In the spring of 2001, almost no one was surprised to hear that several Hollywood studios had been setting up phony web sites to create buzz for new movies. The sites, supposedly run by fans, were just the latest version of some standard tricks in parts of the marketing world.

The exposure of the deception again brought to focus a reality of the modern age: for manipulators, con artists, gossips, and jokesters of all varieties, the Internet is the medium from heaven.

Technology has given us a world in which almost anyone can publish a credible-looking web page. Anyone with a computer or a cell phone can post in online forums. Anyone with a moderate amount of skill with Photoshop or other image-manipulation software can distort reality. Special effects make even videos untrustworthy.

We have a problem here.

CUT AND PASTE, RIGHT AND WRONG

The spread of misinformation isn't always the result of malice. Consider the cut-and-paste problem.

Until recently, people would clip a news article from a paper or magazine. They'd give or mail it to someone else. Now we just copy it digitally and send it along. But when we cut and

paste text, we can run into trouble. Sometimes the cutting removes relevant information. On occasion, words or sentences are changed to utterly distort the meaning. Both practices can prove harmful, but the latter is downright malicious.

In one of the most famous cut-and-paste cases, a column by *Chicago Tribune* columnist Mary Schmich made its way around the Net as a supposed MIT commencement address by novelist Kurt Vonnegut. Schmich had written a wry version of a graduation speech she'd give if asked—"Wear sunscreen," her commencement address began. But somehow, as it spread far and wide, her name came off and Vonnegut's replaced it. (I must have gotten a dozen emails quoting it.) In August 1997, commenting on the case in a subsequent column, Schmich wrote: "But out in the cyberswamp, truth is whatever you say it is, and my simple thoughts on floss and sunscreen were being passed around as Kurt Vonnegut's eternal wisdom. Poor man. He didn't deserve to have his reputation sullied in this way."[242]

Far more troubling was the case of Avi Rubin, a computer scientist and official election judge in the 2004 Maryland primary, who had been fiercely critical of electronic voting machines. He wrote a long article about his 2004 experience with the new machines, and while he maintained his strong objections to flaws in the process, he did make some positive remarks about the machines' potential.[243] His words were then taken out of context, he told me several weeks later, by supporters of the flawed machines. He forwarded me an email from a legislative aide in Ohio that confirmed the misimpression—whether it was inadvertent or deliberate wasn't clear—and he was trying hard to correct it.

I've had material misquoted or misrepresented on a number of occasions. The most telling instance took place in 1997 when I wrote a satiric column—labeled as such—"quoting" an unnamed Microsoft executive admitting to illegal business practices. In the same column, a spokesman for two software-industry trade groups was quoted as admitting his organizations might be making wildly inflated guesses about how much

software is being illegally copied. Finally, I had a spokesman for the PC industry announce the end of the sleazy practice of showing video monitors in computer advertisements, but then, in small print, saying the monitor isn't included.

A week later, after the column had been sent out by the Knight Ridder Tribune wire service, I got a call from an earnest woman at the Business Software Alliance. She was astounded, she said, by the quotes attributed to the spokesman for her organization and the Software Publishers Association. She wanted me to know that no one there could possibly have told me that the software industry was making up its piracy estimates, as my column suggested.

"It was a joke," I said.

There was a pause on the other end of the line. "Oh," she said. It turned out that someone had sent her an email containing the offending quotes, but without the column's introductory line that said, "News stories we're unlikely to read," a missing piece that led to more than one misunderstanding. Indeed, I got a similar call later that day from a well-known public-relations person. She reported that email was flying around Microsoft and her PR firm, with various executives insisting they weren't the unnamed sources in my piece.

It had taken almost no time for the column to morph into an urban legend. Musing about this episode later, I wrote: "Actually, the worst part is that Bill Gates interrupted his speech to world leaders in Switzerland to call and offer me $10 million (plus stock options) to stop writing this column and become the editor of the column he writes for *The New York Times* syndicate. I told my boss and asked for a raise, but for some reason he didn't believe me." Happily, neither did anyone else, this time.

I learned a valuable lesson: email a copy of the entire article, or a URL to the original, and let the reader be the judge. And, as my case suggests, be careful of satire; some people are just too dense to get it.

NEW WAYS TO MISLEAD

In early 2004, John Kerry's presidential campaign drew fire when conservative web critics—and several gullible newspapers—published a composite photograph of him and Jane Fonda, one of the right wing's favorite targets. Kerry and Fonda, in a photo that turned out to have been doctored, were shown "together" at a 1970s rally protesting the Vietnam War.[244] It was unclear who created the fake picture, but the willingness of many people to trust this picture spoke volumes about how easy it is to manipulate public opinion.

Moreover, the incident was only the latest demonstration of a truly pernicious trend of modern fakery. Photos are evidence of nothing in particular.[245] This is why publications that print these kinds of photos are subjected to withering criticism, as was *National Geographic* when it moved one of the Egyptian pyramids in a cover photo. Doctoring photos without clearly labeling them as such is a serious offense in most newspapers and news magazines.[246]

Nothing, in a journalistic sense, justifies blatant deception. But the line between improper doctoring and making an image better is less clear than we might like. For example, simple cropping can remove someone who was in the original picture or it can highlight an important element in the image. Photoshop and other image-manipulation tools give darkroom technicians, who once used various physical techniques to highlight some parts of photos and move others into the background, powerful new ways to alter images.

Even more worrisome is the increasing use of doctored video. It's now common practice for televised sporting events to feature advertising digitally inserted on, for example, stadium walls that are actually blank. The growing field of "product placement"—putting brand-name products into TV shows and movies—is moving closer to the news process, and that should disturb everyone. As the film *Forrest Gump* showed, we can put

someone into a scene who wasn't there in reality; digital technology's steady improvements mean this will become trivially easy.

An element of trickery has been present for years in news programming. For example, the backdrops of urban settings behind anchor people are often inserted electronically. But CBS News, for one, took this to another level in 1999 when Dan Rather's newscast, anchored from Times Square, included digitally created billboards advertising products. At the time, CBS officials said they saw nothing wrong with the practice.[247] This isn't deception on the scale of Jayson Blair, who made up fictitious stories in *The New York Times*, but no responsible news organization should ever insert things into a report that are not really there. If viewers are getting used to this kind of trickery, we're all in trouble.

These techniques are made to order for the Internet, where lies spread quickly and can do enormous damage before the truth catches up. Some of the remedies—including digital watermarking of photos and videos so fakes can be discovered—have surface appeal. But they are not foolproof technically because hackers can consistently defeat such schemes, and they would encourage copyright restrictions even more onerous, and therefore more damaging, to grassroots media and scholarship than the ones currently in place.

WHO'S TALKING, AND WHY?

In 2000, Mark Simeon Jakob put out a phony press release that sent the stock of a company called Emulex into a free fall after credulous news organizations took it seriously. He'd sold the stock short, in effect betting that the price would plummet, and made almost $241,000 before he was caught. He pleaded guilty to a felony and was sentenced to prison.[248]

His offense was egregious. But how much did it differ from chat rooms and discussion boards that have grown so popular in recent years? Pump-and-dump schemers have worked these discussions for years, planting information and then selling or buying accordingly. The Internet bubble was fueled, in no small way, by this kind of behavior—and not just online. Famous Wall Street "analysts" were telling the public to buy shares in companies they were calling dogs in private emails to their colleagues. I have some sympathy for small investors who lost big in the bubble, and contempt for the people who knowingly touted absurdly overpriced stocks. But greed was everywhere, and small investors who were looking for something that was too good to be true violated common sense.

Yet the investment forums can be a source of incredibly good information, too. Sometimes disgruntled employees post insider tales that can be a warning of harder times to come for shareholders. Sometimes a particularly bright amateur analyst spots something relevant the pros have missed. To dismiss all online information out of hand is as foolish as ignoring it entirely—but the failure to do one's homework before making a serious decision may be the most foolish mistake of all.

In doing homework, one of the most crucial exercises is to consider the source. Good journalists know this as a matter of practice. We don't pick a random bystander and assume he's an expert on, say, nuclear power. And we'd laugh out loud at the notion of reading some anonymous Net posting and using it as the factual basis for an article—at least I would.

Internet gossip monger Matt Drudge doesn't practice what I'd call respectable journalism (and, to be fair, he doesn't call himself a journalist), but I respect him for this much: he signs his name to everything he posts. That probably didn't come as much consolation to John Kerry, the 2004 Democratic presidential candidate. Kerry, you may recall, was dogged in early February by a rumor of an extramarital affair, a "scandal"—for which there was absolutely no evidence and which was flatly

denied by everyone supposedly involved—that got its legs after Drudge published it on his web site.[249]

Unfortunate as the entire "Kerry affair" may have been, at least we knew who was largely responsible for having put it into play in the first place. And we could weigh the allegations in the context of the writer's previous work. However, we can't make such judgments about a lot of other things we read online. One of the Net's great features, the ability to remain anonymous, can also be one of its chief defects.

People I respect have told me we need to do away with anonymity on the Net. They have good reasons.

But anonymity is enshrined in our culture, even if its use can be distasteful at times. And there are excellent reasons for keeping one's identity hidden. A person with AIDS or another disease can lose a job or housing, or be persecuted in more violent ways. Someone holding unpopular political views in a small town that leans strongly in one direction may want to discuss it with others of like mind. Corporate and government whistle blowers need to be able to contact authorities and journalists without fear of being revealed. More than anyone, political dissidents in nations where such behavior can be life-threatening deserve the protection of anonymity when they need it.

Though the benefits of anonymity are clear, it also has its hazards. In one now famous example in 2004, a software glitch at Amazon.com revealed what many people suspected about the site's customer-written book reviews: authors were penning rave reviews of their own work under false names and, in some cases, slamming competing books. A *New York Times* story[250] showed a remarkable willingness on authors' part to excuse their deceptions as just another marketing tool. A more reasonable excuse was counteracting trash reviews by enemies. I worry what will happen when this book is published. I certainly have my share of adversaries. Will they trash me on Amazon? No doubt. Will that hurt sales? Probably. Can I do anything about it, assuming they don't libel me? Probably not.

In one online discussion on my blog about copyright, I challenged a commenter named "George" on his refusal to say who he was. "You're welcome to remain anonymous," I said. "I think you would enjoy even more credibility in this discussion if you said who you were. A casual reader might wonder why you want to be anonymous."

He replied: "You should judge my credibility by how my statements correspond with the facts, logic, and the law—not by who I am."[251]

He had it partly right. Debating skills are not proof of anything. In the absence of a foundation for his comments, he hadn't earned anyone's trust. Credibility stems not just from smart arguments; it also comes from a willingness to stand behind those arguments when a compelling reason to stay anonymous is absent. There was none in this case.

Another commenter, also using a false name, defended an electronic voting machine maker's use of copyright law to suppress memos that revealed flaws in its voting systems. It seemed that he or she was also posting comments, using a different name but similar (and in some cases identical) language, on a blog about intellectual property sponsored by the University of California-Berkeley journalism school. I learned this because Mary Hodder, one of the principal authors of that blog,[252] noted similarities in style in postings on our respective sites, which we believe share a number of readers due to the topics we cover. We checked the Internet addresses from which the comments had been posted; they were identical. This didn't absolutely prove that the same person was making both comments, but it helped make the case. Not only was this person refusing to be identified, but he or she was trying to make it seem as though a posse was patrolling our blogs to show us the error of our ways when, in fact, it was just one person on both.

What do these examples suggest? People reading comments on discussion boards would be wise to question the veracity of a commenter whenever they aren't absolutely sure where the posting is coming from.[253]

As we discussed in Chapter 8, advances in technology are likely to bring us better ways to gauge and, in effect, manage reputations and verify a commenter's bona fides without exposing his or her actual identity to the world.

Googling someone, to see what else he or she has said online in other places, sounds like a good way to start. But it ultimately isn't the answer. If, however, someone has been using a consistent pseudonym, at least we have the possibility of knowing if a person is reputable or has been making trouble elsewhere.

At the moment, my favorite solution is not the most practical: if everyone had a blog or other kind of web site, they could include a link as a kind of digital signature. Yes, web sites can be faked, but a hoax that uses someone else's name or hides behind a pseudonym for improper purposes, could attract unwelcome attention from the authorities—and because web site owners have to pay someone for hosting their site, the owner can be traced. Again, I would do nothing to stop anonymity on the Internet. But if we are going to have serious online discussions, I think all parties should, with few exceptions, either be willing to verify who they are, or risk having their contributions be questioned and, in some cases, ignored.

TROLLS AND OTHER ANNOYANCES

Grassroots journalism has more problems than deciding whether anonymous posting is a good or bad idea. For starters, consider the trolls.

Rob Malda, Jeff Bates, and their colleagues at Slashdot have been dealing with trolls for years. At Slashdot, subtitled "News for Nerds: Stuff that Matters," the readers do the heavy lifting. They're constantly combing the Web for interesting information—articles, news stories, press releases, and mailing list postings—and recommend the material to Slashdot's tiny

editorial staff. Each day, the editors select a dozen or so of the best items, which they highlight on the Slashdot homepage with a short summary and hyperlink, and invite readers to comment online. Then the editors sit back to watch what happens, and so do hundreds of thousands of other people.

The initial summaries and links are the beginning of the conversation on Slashdot, not the end. The average item generates about 250 comments. Some generate far more. Moderators, themselves selected on the basis of their participation in other discussions, rate the quality of the postings, and readers can adjust the results so they see everything or, as most do, a subset of the more substantive comments.

The Slashdot team has had to keep tweaking the software that runs the Slashdot site, as well as the user-based moderation system, because of the trolls and vandals who try to clog the site with irrelevant or obscene postings, ruining the experience for others. It's a constant annoyance, Bates told me, but part of the price of doing business.

How do you know if a troll is on your site? The definition on Ward Cunningham's Wiki says it best:

> A troll is deliberately crafted to provoke others with the intention of wasting their time and energy. A troll is a time thief. To troll is to steal from people. That is what makes trolling heinous.
>
> Trolls can be identified by their disengagement from a conversation or argument. They do not believe what they say, but merely say it for effect.
>
> Trolls are motivated by a desire for attention by people and can't or won't acquire it in a productive manner.
>
> Someone may be insufferable, infuriating, fanatical, and an ignorant idiot to boot without being a troll.
>
> Also note that a troll isn't necessarily insulting, snide, or even impolite. Only the crudest, most obvious, forms of trolling can be identified so easily.
>
> If you find yourself patiently explaining, at length and in great detail, some obscure point to someone who isn't even being polite to you, then you are probably being trolled.[254]

User registration on comment systems, with a name and verifiable email address, can be a deterrent to trolls. The worst thing you can do, as Netizens know, is feed the troll. Ignoring him is usually the best answer. If people become abusive, they can be banned from discussions. Not everyone has a right to speak on everyone else's site or be part of everyone else's conversation.

SPIN PATROL

Journalists become accustomed to a process known as spinning. Wikipedia accurately describes this, in the context of public relations, as "putting events or other facts, especially of those with political or legal significance, into contexts favoring oneself or one's client or cause, at least in comparison to opponents. Newmakers and their PR legions have been spinning us since the media became a way to get information to the public, and we've been alternately falling for it or resisting it all this time."

In the physical world, I always try to ask myself what a person I'm interviewing has to gain from doing an interview. We need to recognize that motives play a part in what we're told, and we adjust our ultimate coverage accordingly.

But spin takes some insidious routes to the public. One of the worst forms is the media's lazy use of press releases as news. Some smaller newspapers are known to print them verbatim, as if a reporter had actually done some reporting and writing. Lately, video press releases have become a stain on both the PR profession and journalism. Local TV stations are handed video releases, often including fake "reporters" interviewing officials from the company or government agency that wants to get its news out, and too often stations play all or part of these mockeries of journalism. In March 2004, the Bush administration

was properly chastised for sending out video releases to promote, in a highly political way, a drug-benefits bill Congress had passed a few months earlier.[255]

Online spin varies from the relatively harmless, and even amusing, to more ethically challenged methods. On the harmless side is "Google bombing," a method of connecting a word or phrase to a specific web site through the Google search engine. After one group of Google bombers got "miserable failure" to point to George W. Bush's biography page on the White House site, his supporters retaliated by connecting John Kerry's page to the word "waffles."[256] Sooner or later, Google will either prevent this kind of thing or risk some of its own credibility.

Cyber-spin is getting more sophisticated, especially when it comes in comments or other postings by someone who's trying to make a point but doesn't identify his or her connection to the subject. The entertainment-industry copyright defender who made such a point of critiquing my blog was, in effect, spinning not just me but my audience as well. This is an unintended effect of the conversation, but one we'll have to live with.

Just before the January 2004 Consumer Electronics Show, I got an email from someone telling me, in a fairly breathless way, about a product due to be announced at the show. He was gleeful, it seemed, that the company had inadvertently given out information it intended to keep under wraps until the official announcement. He pointed me to several pages, including one that had a picture of the gadget (some gear for networking multimedia at home) and another where the company's chief executive had essentially confirmed the product's existence on a product support forum.

So I posted this information on my blog. "Consider this a small example of tomorrow's journalism today," I wrote. "A reader who knew much more than I did about something did some reporting and found information worth noting. Now you know, too."

Was I spun? After all, it wasn't a product I was likely to cover in the first place. My guess, based on some follow-up checking, is that this wasn't spin but a tip from someone who really thought he was giving me a scoop. Still, I plan to be more cautious before posting such things in the future.

Some online spin is obviously deceptive, as Adam Gaffin discovered. Gaffin runs an online forum called "Wicked Good" on the Boston Online site.[257] A 2003 forum thread talked about a fictional company in a soap opera holding a "Sexiest Man" contest. Someone named "dixie wrecked" was talking up the contest and the TV show. Gaffin got suspicious and checked the Internet address from which "dixie" was posting, and discovered it originated at a Washington-based firm, New Media Strategies, a company that offers, according to its web site, online word-of-mouth marketing to create buzz about products and brands. "We've been played," Gaffin told his forum[258] members, adding, "So, just in case Google indexes this page: New Media Strategies sucks. Let me repeat, New Media Strategies sucks."

Interestingly, by early 2004, one item on the first page of Google listings using the search term "new media strategies" was a pointer to a Boston Online page entitled, "Why New Media Strategies sucks." (The item had moved down to the second page by late April.)

I don't mean to pick on New Media Strategies here, or to suggest that its mistake in this case represents the company's general methods.[259] I do want to suggest that just one such episode, if it's caught and then stirs up any degree of irritation online, can be a lasting blemish.

Another lesson: exposure can be a reasonable counterweight to spinmeisters. Unfortunately, not everyone can catch such acts. We need better ways to sniff them out and then expose them with a variety of tools, including reputation systems. In many cases, the best solution is to ensure an open conversation among informed readers because they'll collectively inform each other.

CITIZEN REPORTERS TO THE RESCUE

Blogger Ken Layne[260] captured one of the online world's essential characteristics in a classic posting in 2001. "We can Fact Check your ass," Layne said.[261] When there are lots of citizen reporters scrutinizing what other people say, they have a way of getting to the truth, or at least shining light on inconsistencies.

Case in point: Kaycee Nicole created a blog to talk openly about life, illness, and loss. As she grew sick and lay dying, she created a community. Thousands of people visited her blog in 2000 and 2001. They comforted her—and each other—with messages of support and offers of help. They researched her illness, looking for a way to make her better. And Kaycee did get better, at least for a while. Then she sickened again and finally succumbed to her leukemia.

But on May 18, 2001, someone named "acridrabbit" posted a simple question on MetaFilter, a collaborative blog and news site: "Is it possible that Kaycee did not exist?" The query set off a furious controversy. A relatively small but relentless group of Net denizens unraveled the tale of anguish and discovered a hoax. They investigated court records. They checked their findings with each other. They did some of the best detective work you'll ever see.

What this group accomplished was, in a sense, investigative reporting. But they weren't professional journalists. They were strangers who, for the most part, only knew each other online. But combining the power of the Internet and old-fashioned reporting, they'd come together—first in sorrow, then in dismay that morphed toward anger—to scrutinize a situation and, ultimately, solve a mystery.[262]

Fact-checking is a just one tool a community can bring to bear. As in open source projects, combining all those eyes and ideas can create a self-righting phenomenon. In the summer of 2003, David Weinberger and I discovered other community benefits. We'd launched a small, noncommercial web site called

WordPirates,[263] the purpose of which was to remind people how some good words in our language have been hijacked by corporate and political interests.

We opened the site to allow anyone to add a word plus an explanation of why it should be there. As we expected, some folks used the system to make off-point, irrelevant, or puerile postings, often with no explanation. We've had to prune heavily.

But a vandal found a security flaw in the software powering the site and exploited it by posting programming code inside a comment form—some HTML that took users to an unaffiliated web page containing one of the most disgusting photographs I've ever seen. We removed the offending post, thanks to a sharp-eyed programmer who let us know how the page had been misused so foully. Finally, the developer of the software we were using, who hadn't anticipated this kind of abuse, fixed the security breach.

We'd surely seen the downside of the Net. But we also saw the upside in the way the community helped us find, analyze, and fix the problem. As Weinberger noted after our dust up with the rogue coder: "It's as if the Internet is not only self-correcting about matters of fact but also morally self-correcting: A bad turn is corrected by several good ones."

A FLIGHT TO QUALITY?

The flood of unreliable information on the Net could have the ironic effect of reinforcing the influence of Big Media, at least in the short term. This assumes, of course, that users of online journalism trust Big Media in the first place. Many do not.

Unlike many Americans, and in spite of some media scandals, I have substantial faith that major newspapers try hard to be accurate and fair. For example, I've been reading *The Wall Street Journal* for years, and I trust that the typical front-page

news article in the *Journal* has been well reported, written, and edited. That doesn't mean I assume that everything in it is true, though I do assume the paper has done its best, and that there are institutional mechanisms in place to correct something if it's wrong. Those beliefs have carried over as, increasingly, I read the *Journal* online rather than in print. (Even after the Jayson Blair mess, I'd say the same thing about *The New York Times*.)

But Big Media, as it participates in the new conversation online, takes on risks that could hurt credibility even more. One of these days, someone is going to break through the security of a major media web site—the *Journal* or the *Times* or CNN—and post some "news" that turns out to be absolutely false. Maybe the story will announce wonderful news for some company, or terrible tidings, thereby giving the unscrupulous computer crackers, terrorists, or even politically connected malefactors a way to manipulate the stock market, cause panic, or steal an election.[264]

This act, which I consider more a certainty than a possibility, will change the news media's trust equation, at least for a time. Will it have long-lasting impact? Only if it happens repeatedly.

PLAIN OLD COMMON SENSE

Being a reporter involves some basic practices. When I see or hear about something I think may be worth reporting to my audience, I verify it, or quote credible people who should know, or go to the source (human or document). If I link to something intriguing on my blog but don't know whether it's true, I offer that caveat. Generally, I don't just repeat an anonymous posting. If the fact in question didn't come from a source I trust, I check it out.

Users of online information need to develop similar filters. They need a hierarchy of trust.

In my own hierarchy, I trust *The New York Times* more than a supermarket tabloid. I trust what Doc Searls tells me on his blog more than what a random blogger says on a page I've never seen before.

As noted earlier, we need better recommendation and reputation tools, software that lets us traverse the Web using recommendations from trusted friends and friends of friends. We'll be figuring this out in the next few years, and I'm confident we'll get better and better at it.

But for now, people need to take information on the Internet with the proverbial grain of salt. When they see things that promise a measurable impact on their lives—such as a news story that persuades them to sell or buy something expensive—they should verify the claim before reacting.

There are limits to this, but on matters where the personal stakes are sufficiently high, it's probably worth remembering the legendary admonition given by crusty old editors to green reporters: if your mother says she loves you, check it out.