

1

Comment: The Bottom Half of the Web

There's a reason that comments are typically put on the bottom half of the Internet.

—@AvoidComments (Shane Liesegang), Twitter

“Am I ugly?” This question has been asked on YouTube by dozens of young people, and hundreds of thousands of comments, ranging from supportive to insulting, have been left in response. The Web is perfect for this sort of thing, and it has been almost from the start—even if some think it alarming. Over a decade ago, some of the earliest popular exposure that the Web received was through photo-rating sites like HOTorNOT. “Am I Ugly?” videos continue this phenomenon and remain true to YouTube’s origins. YouTube was conceived, in part, as a video version of HOTorNOT. YouTube cofounder Jawed Karim was impressed with the site “because it was the first time that someone had designed a Website where anyone could upload content that everyone else could view.”¹ But not only could people upload content: others could comment on that sophomoric content. And if the word *sophomoric* seems haughty, Mark Zuckerberg was a Harvard sophomore when he first launched Facemash, his hot-or-not site that used purloined student photos from dormitory directories, what Harvard calls “facebook.”

This uploading and evaluating of content by users is now associated with various theories and buzzwords. (I return to the question of others’ physical attractiveness in a later chapter.) *Social media*, like YouTube, are populated by *user-generated content*. Facebook is an example of *Web 2.0* in that it harnesses the power of human networks. The online activity of masses of ordinary people might display the *wisdom of the crowd* or *collective intelligence*. Books on these topics claim that this “changes everything” and is transforming the Internet, markets, freedom, and the world. Yet I continue to be intrigued by what is

happening in the margins—the seemingly modest comment. But what is comment?

As I use the term, *comment* is a genre of communication. In 2010, for instance, YouTube lifted its fifteen-minute limit on videos, and since then, there has been a flurry of ten-hour compilations of geeky, catchy, and annoying audio. Under a hypnotic video of Darth Vader breathing, someone commented: “What am I doing with my life?! 10 hours of breathing!”² From this we can see that comment is *communication*, it is social, it is meant to be seen by others, and it is *reactive*: it follows or is in response to something and appears below a post on a blog, a book description on Amazon, or a video on YouTube. Although comment is reactive, it is not always responsive or substantively engaging. Many comments on social news sites are prefaced with the acronym *tl;dr* (too long; didn’t read), meaning that the commenter is reacting to a headline or blurb without having read the article. Comment is *short*—often as simple as the click of a button, sometimes measured in characters, but rarely more than a handful of paragraphs. And it is *asynchronous*, meaning that it can be made within seconds, hours, or even days of its provocation. Putting aside future transformations, comment is already present: comment has a long history (some of which I discuss briefly), and it is pervasive. Our world is permeated by comment, and we are the source of its judgment and the object of its scrutiny. There is little novelty in the form of comment itself, but its contemporary ubiquity makes it worthy of careful consideration, especially given online comment’s tarnished reputation as something best avoided.

This understanding of comment as communication that is reactive, short, and asynchronous fails to draw a bright line. (I use the term *comment* to speak of the genre and reserve *comments* for an actual plurality of messages.) For instance, at what point does a message become too long to be considered a comment? Unlike a tweet, there is no character limit for a comment, but I focus on communication that is relatively short and can live outside the expectations of real-time interaction. And although these are the rough contours of comment, its essence is best expressed by way of—appropriately enough—an online exhortation: “Don’t read the comments.” This popular maxim is captured in the tweets of game designer Shane Liesegang. At his account @AvoidComments, he claimed that “there’s a reason that comments are typically put on the bottom half of the Internet.”³ There is a lot of dreck down there, but in sifting through the comments, we can learn much about ourselves and the ways that other people seek to exploit the value of our social selves. This book is an exercise in reading

(rather than avoiding) comment, and it documents an expedition to the bottom of the Web. I show how comment can inform (via reviews), improve (via feedback), manipulate (via fakes), alienate (via hate), shape (via social comparison), and perplex us. I touch on the historical antecedents of online comment and visit the communities of Amazon reviewers, fan fiction authors, online learners, scammers, free thinkers, and mean kids.

The point of this journey is expressed by an adage often used by media theorist Marshal McLuhan: “We live invested in an electric information environment that is quite as imperceptible to us as water is to fish.”⁴ Comment is easily seen but invisible to the extent that we take it for granted. Often when comment does make an impression on us, it is a nuisance to be disabled or an offense to be ignored. And even when we do see and appreciate comment, most people have no idea the extent to which it is manipulated. For example, people *like* things on Facebook over four billion times a day, and even these littlest of comments are big business. Scammers are proliferating “*like farms*.” When purported pages about cute puppies, brave veterans, and young people stricken with cancer have enough likes, their content is decorated with ads and links to malware sites, or they are sold to others who will do the same.⁵ It is easy to appreciate why some recommend that we “never read the comments.” Much like California during its gold rush, the bottom half of the Web can be lively and lawless, and it is where many are attempting to make a fortune. Although I do not advocate that everyone read all the comments all the time, I think that it is wise to understand them.

The easiest way to avoid comments is not to have them. Because many sites have disabled their comments, I begin this journey with what gossip teaches about online discussion and why many users and sites are turning away from comment. I argue that disabling comment is a reflection of a platform’s growth as users seek intimate serendipity and flee “filtered sludge.”

Gossip

The origins of YouTube and Facebook demonstrate that people like to talk about one another: we gossip. Although gossip might seem like a trivial thing, evolutionary psychologist Robin Dunbar argues that it is central to understanding humanity. If you participate in online communities, you might have heard of Dunbar’s eponymous number of 150. People invoke *Dunbar’s number* when a

community (such as an email list whose members used to know nearly everyone else on the list) grows too big. The Web is a big place, and any technology that permits its denizens to communicate with one another has to grapple with the problem of social scale. After the group becomes too large, people complain that the magic has gone. Graffiti and scams proliferate. The known personalities and easy cadence of the group have been replaced by strangers and bickering about unruliness and the need for moderation.

Dunbar did not set out to coin an aphorism about online community. He was seeking to answer the question of why primates, especially humans, are smart—why human brains are about nine times larger, relative to body size, than the brains of other animals. Some have suggested that brain size was related to environment, the use of color vision to find fruit, distances traveled while foraging, or the complex omnivore diet. When Dunbar looked at all of these variables among primates, however, he found no such pattern. But the size of primates' neocortex did correlate with the size of their groups and the time that they spent grooming one another.⁶

A large group is better protected against predation than a small group is, but it also has internal competition for food and mating. Even monkeys can scheme and are sensitive to threats from their peers. Grooming, therefore, is an activity through which alliances are forged and disputes resolved. Experiments with wild vervet monkeys show that they are more likely to pay attention to the distress calls of individuals with whom they recently groomed. But keeping track of who is scratching whose back can be complicated. In a group of twenty, there are nineteen direct relationships and 171 third-party relationships, so as group size increases, so does the complexity of the network and the time that primates spend grooming one another. According to Dunbar, primates' large brains are the result of an evolutionary race of alliances through social grooming.

For humans, social grooming includes language. Because larger groups require more efficient means of forging alliances, gossip circulates information about others in the social networks in which they exist. While the practice of talking about others (including rumors and bathroom graffiti, or latrinalia) is more interesting and complex than I can address here, I understand gossip simply as “evaluative social chat.”⁷ And the alliances that result from sharing opinions about others can be Machiavellian. On a television reality show, for example, Sandy might realize that John's seeming betrayal of Alice could itself be a lie. Dunbar argues that gossip requires a sophisticated type of social cognition known

as the *theory of mind* through which we infer the mental states of others. Even four-year-old children demonstrate second-order intentionality: the child has a belief about what someone else wants. Adults can negotiate fourth- or even fifth-order intentionality. This is amusingly demonstrated in a scene in *The Princess Bride* where two opponents engage in a battle of wits. The “Man in Black” poisons one of two goblets of wine, and Vizzini must choose and drink from the safe one. Vizzini begins his chain of deduction with the assumption that “A clever man would put the poison into his own goblet, because he would know that only a great fool would reach for what he was given. I am not a great fool, so I can clearly not choose the wine in front of you. But you must have known I was not a great fool. You would have counted on it, so I can clearly not choose the wine in front of me.” This type of inference requires a “clever man” with a big brain.

In any case, Dunbar’s number of 150 is, roughly, the cognitive limit of how many relationships humans can maintain given their complexity (such as “the enemy of my enemy is my friend”). However, Dunbar proposed multiple tiers, from the family up to the tribe of a couple thousand people. The rough size of the *clan*—the individuals that a person keeps in contact with and can track relationships for—is 150 people. This is roughly the size of early farming communities, modern planters in Indonesia and Philippines, and contemporary Hutterites. It is roughly what the Church of England concluded to be the ideal size for a congregation and the number of soldiers in a military company. Also, using the birth rates observed in hunter-gatherer or peasant societies, this number corresponds to about five generations, which is as far back in time as anyone living can remember: “only within the circle of individuals defined by those relationships can you specify who is whose cousin, and who is merely an acquaintance.”⁸

How is Dunbar’s number related to comment? It provides an unexpected clue to why comment frequently fails on the Web.

Disabling Comments

In the blogging domain, there is little that Dave Winer has not written code for, started a company around, or opined about. He often is credited with deploying the first blog comments in 1998. (Another contender for this claim is Bruce Ableson at Open Diary.⁹) Despite his quick smile and easy manner, Winer ends up in a lot of online arguments. In 2001, he was described in the *New York Times*

as someone who is “not shy about ruffling the big names in high technology.”¹⁰ Winer is opinionated and passionate. He also is willing to say that something “sucks” or to call someone an idiot. As I will discuss further, *drama genres* of comment, such as sites where people can ask others questions or make lists of things to avoid, lend themselves to conflict. Certain personalities also seemingly attract conflict. In 2006, when Winer announced that he would stop blogging before the end of the year, a group of antagonists created a countdown clock that they said would continue “until he shuts up.” (He continued blogging.) Critics expressed their antagonism in comments to his blog, and he repeatedly considered disabling the comments altogether. In 2007, he wrote that he did not think comments were an essential part of a blog, especially if they “interfere with the natural expression of the unedited voice of an individual”:

We already had mail lists before we had blogs. The whole notion that blogs should evolve to become mail lists seems to waste the blogs. Comments are very much mail-list-like things. A few voices can drown out all others. The cool thing about blogs is that while they may be quiet, and it may be hard to find what you’re looking for, at least you can say what you think without being shouted down. This makes it possible for unpopular ideas to be expressed.¹¹

In 2010, Winer developed this idea further, arguing that blog comments should be short and about the blog posting (or *responsive*, using my term). They ought not be digressive or overly long. He proposed a system that allows comments of less than a thousand characters to be invisibly submitted within twenty-four hours: “After the commenting period is over, the comments would become visible, and no further comments would be permitted.” Those who wished to respond later or at greater length could do so on their own blogs. This idea was supported by the *trackback* feature of many blogs. If I respond to a blog post by Winer with a post on my own blog, for example, then my blogging platform would inform Winer’s blog service of my response. Winer’s blog entry would then include a link to my own. Winer’s blog “tracks back” to those who respond. Trackbacks were seen as a way to complement or replace comments, but they have largely fallen into disuse after their abuse by spammers. For Winer, neither blog comments nor tweets are appropriate for conversation. In 2012,

Winer, the person who often is credited with first enabling blog comments, disabled them from his own blog — seemingly forever.¹²

[figure 1.1 near here]

As Mathew Wingram, a technology writer, noted about a 2007 fracas over Winer’s “Why Facebook Sucks” posting, Winer’s approach sometimes “brings the hate.”¹³ But Winer’s experience is not unusual. After the halcyon days of blogging, many bloggers abandoned their sites or shuttered their comments. Some popular sites (including *Boing Boing* in 2003, the *Washington Post* in 2006, *Engadget* in 2010, and *Popular Science* in 2013) have turned off their comments for extended periods. In 2013, Rob Beschizza, managing editor of *Boing Boing*, tweeted that he may do so again “for good,” perhaps after inappropriate comments were made on a posting about the death of a friend.¹⁴ *Boing Boing* began as a paper zine in the late 1980s and went online in 1995. In its early days, it was like an informational swap meet among friends, but by the new millennium, it had become too popular to serve as an unfettered venue for sharing and gossiping among friends and has struggled with this fact ever since.

However, there are two other responses to unruly comments beyond disabling or ignoring them. Website managers can attempt to *fortify* their commenting system and Website users can *relocate* in search of what I call *intimate serendipity*.

Fortifying Comment Systems

By *fortify*, I mean to make the system more resistant to abuse. Some sites require users to perform a task (like typing in text that has been distorted) before leaving a comment so as to minimize abuse. However, abusers often match the cleverness of the challenge or farm out the task to low-cost workers on the other side of the world. Many sites permit readers to filter comments based on ratings from other users who act as moderators, such as at the nerdy news site Slashdot. This site also uses meta-moderation, whereby others’ moderations can be rated as fair or unfair. Even so, people sometimes complain that a *cabal* of moderators has taken over and is abusing the system. That is, a group of users collude to promote one another’s postings and standing. One can often find someone complaining in

a story's comments that their submission and summary of the story was better and earlier but that it was ignored because she was not part of a clique.

Facebook and Google+ have required users to use their real names. While Facebook has been relative lax in enforcement, Google+ was quite strict at its start but stepped back from the requirement in 2014. (My dog has a Facebook page, but it is under his real name.) Such social networks are then able to leverage their identity policies and reach by providing authentication and commenting services for others. *Slate* adopted Facebook's 2011 "Comments Box" service, and Farhad Manjoo, a staff writer at *Slate*, was pleased that Facebook knew real names and that comments could be seen on Facebook by the commenter's friends and family: "This introduces to the Web one of the most important offline rules for etiquette: Don't say anything that you'd be ashamed to say in front of your mom."¹⁵ MG Siegler, a blogger at *TechCrunch*, noticed that since they "flipped the switch" and adopted Facebook's service there had been a large drop in comments, but "this is completely expected and definitely not a bad thing." Before, a post might get hundreds of comments, half of which were "weak to poor" and half of those "pure trollish nonsense." With the new system, a similar post might receive about a hundred comments, half of which "are actually coherent thoughts in response to the post itself—you know, what a comment is supposed to be."¹⁶ Others claimed that real-name policies suppressed anonymous speech and were incompatible with the multiple identities that we maintain in life. Furthermore, using centralized commenting systems cedes ever more autonomy and privacy to the likes of Google, Facebook, and other comment-specific services like Disqus (used by 750,000 sites, including CNN's Website) and Livefyre (used by the BBC and the *New York Times*).

Sometimes, relatively simple approaches do the trick. The link-sharing and discussion site MetaFilter requires a one-time \$5 membership fee. This fee, its strong community norms, and occasional moderation for flagrant abuse seem to have fostered a robust and civil culture. Newspapers, too, have experimented with asking readers to subscribe or pay a small fee to comment. In keeping with Dunbar's insight, Clay Shirky, author and NYU professor, is fond of examples of communities whose size is purposefully limited. Ten years ago, he wrote about an email list that removed the oldest subscribed member to make room for the newest one. Another list's periodic purge was inspired by a supposed neighborhood hot tub that was accessed by a key-coded gate lock: people could give the code to friends, but when the bathers became too rowdy or the tub was

overcrowded, the owner simply changed the code and gave the new one to his immediate friends under the same policy.

A decade later, Shirky continued to stress that “intimacy doesn’t scale.” In a 2013 talk entitled “Why Do Comments Suck?” he put it simply: “Comment systems can be good, big, cheap—pick two.”¹⁷ Many sites with comments seek a large audience. They “want lots of people to forward the article to a million friends, shut up and then read another article.” Sites that treat their users as community members (through smaller size or careful moderation) tend to have better comments. This is a good insight but easier said than done: good communities tend to grow. This is the paradox of their success. People then often relocate to another site without a good understanding of what went wrong (except that it went “downhill”) or what they were looking for in the first place.

Twitter and the Search for Intimate Serendipity

We now have a toolbox of tactics for resisting comment abuse, and they often are good enough—for a time. But some communities struggle as they experiment with finding a system that is appropriate to their changing size and circumstances. Those not satisfied with the changes often leave and *relocate* to a new media platform in search of what I call *intimate serendipity*. When I went to blogging get-togethers in 2003, it was with a dozen of like-minded enthusiasts: I met interesting people and we had good conversations. Over a decade later, going to a meeting for people who post comments to the Web seems passé. (Today almost any gathering could qualify as such a meeting.) After a network of people (online or otherwise) becomes popular, people want to bring their friends. At first, this is great. The value of a network increases significantly with each new node. A network of five phones permits ten connections; doubling the phones to ten permits forty-five possible connections. As Dunbar notes, however, at some point the scale of networks overwhelms the participants. First, we ask, “Who brought that guy to the party?” Second, the network becomes a target for those who wish to exploit it via spam and manipulation.

I sometimes ask my students if the parade of Web platforms is over (i.e., Geocities, Myspace, Facebook, Twitter, and Google+). Given the difficulties involved in leaving a service (because content and connections must be abandoned) and the profusion of niche networks, some might think that there is little need for anything new. But people do relocate when an existing platform

becomes overly populated by jerks, spammers, and ads or overly constrained by controls and filters. People often want a network where intimacy and serendipity are possible. Although there are sites where being anonymous and a jerk is the norm, many people want to express their authentic selves without fearing attacks, manipulation, or unusual exposure and while remaining open to things that surprise and delight. When they can't do so, they move, as seen in the migrations of geeks between social news sites like Slashdot, Digg, Reddit, and Hacker News.

[figure 1.2 near here]

By 2007, three years after its founding, Digg was being criticized as a system that was rife with manipulation by supposed “bury brigades” that suppressed others’ stories. At the same time, the company was trying to become financially viable or even profitable. As users began to leave, the site exacerbated discontent with the introduction of unpopular changes, including new comment systems and site designs. By 2010, the site was on its deathbed; while the service limped on, and the name had been affixed to various “reboots,” its founder, staff, and community were gone. Many who left Digg relocated to Reddit. Correspondingly, those who had been at Reddit, especially the technically inclined, lamented that the site had become too big and diluted. When Paul Graham launched Hacker News in 2007, the intention was to replicate the early Reddit days, and some early contributors to Reddit followed.

The irony is that success sometimes, seemingly, brings about a comment system’s demise. In 2009, law professor and civic reformer Lawrence Lessig announced that he was retiring his blog. Lessig’s problem was not with haters but with the love. He is an influential legal scholar, a successful author, and a founder of significant cultural, academic, and civic organizations. He has argued against copyright extensions before the U.S. Supreme Court, is a founder of Creative Commons, and popularized the notion of “free culture” with a successful book. However, a growing family, a shift in research focus, and the “increasingly technical burden to maintaining a blog” were too much. He and his volunteers tried to keep up, but a third of the thirty thousand comments on his blog were likely “fraudsters,” and online casino spam was growing. Google stopped indexing the site at one point. However, he was “still trying to understand

twitter.”¹⁸ In fact, lots of people were moving to Twitter. For a time, it gave its users intimate serendipity.

At its beginning in 2006, Twitter felt edgy and intimate. It was not uncommon to find users flush with an encounter with a (minor or major) celebrity.¹⁹ Also, people (especially the famous) were thrilled to be able to authentically express themselves. People like talking about themselves, and Twitter appeared to be a safe space to do so. Research indicates that people spend 30 to 40 percent of their interactions telling others about their subjective experiences, so it is not surprising that researchers found that 41 percent of the tweets in their study were of the “me now” variety.²⁰ Two Harvard University neuroscientists concluded that “disclosing information about the self is intrinsically rewarding” because it triggers regions of the brain that are associated with the mesolimbic dopamine reward system. In experiments in which subjects could choose to speak about themselves or factual matters, people chose to speak about themselves the majority of the time. When these choices were associated with small payments, people were willing to pay an average 17 percent “tax” to talk about themselves: “Just as monkeys are willing to forgo juice rewards to view dominant groupmates and college students are willing to give up money to view attractive members of the opposite sex, our participants were willing to forgo money to think and talk about themselves.”²¹

Trent Reznor, the digitally progressive artist behind the band Nine Inch Nails, initially took advantage of this self-disclosure. At first, Twitter allowed him to “lower the curtain a bit and let you see more of my personality,” more so than what he could do on the forums at his site nin.com:

The problem with really getting engaged in a community is getting through the clutter and noise. In a closed environment like nin.com a lot of this can be moderated away, or code can be implemented to make it more difficult for troublemakers to persist. It’s tedious and feels like wasted energy doing that shit, but some people exist to ruin it for others—and they are the ones who have nothing better to do with their time. Example: on nin.com, there’s 3–4 different people that each send me between 50–100 message per day of delusional, often threatening nonsense. We can delete them, but they just sign back up and start again. Yes, we are implementing

several changes to address this, but the point is it quickly gets very old weeding through that stuff.²²

At Twitter, Reznor “approached that as a place to be less formal and more off-the-cuff, honest and ‘human.’” Some of his tweets were about a new love in his life, which some fans found incompatible with his earlier dark and alienated music: “I’m not the same person I was in 1994 (and I’m happy about that).” This worsened in 2009, when the obsessively vitriolic fans he called “the Metal Sludge contingency” discovered Twitter; he vowed to quit and has only tweeted sporadically since. Unlike email, where addresses can be kept private and messages can be filtered, a Twitter handle is public. Such openness worked as long as the community was sufficiently small, but by 2009, the tweet bomb had arrived: people were spamming others. Twitter now permits users to block or label others as spammers, and serious users use third-party apps with powerful filters. But new Twitter accounts are easily created or cheaply purchased by the thousands, as I discuss in a later chapter. By summer 2013, people were using a “block bot” to filter out “general bigots, assholes and fools.”²³ Yet, a centralized block list is controversial for many reasons, including who decides who goes on the list.

Many social platforms move from intimate serendipity toward filtered sludge, and some manage it better than others. As investors begin to demand a return on their investment, the sites themselves are tempted to alienate their users with ever more intrusive filtering and ads. This is the life cycle of a social media platform. Although I expected Twitter to face the same dynamic, I was surprised that its crisis began with a case of humanitarian advocacy.

The Tweet Bomb and the Loss of Intimate Serendipity

In spring 2012, I was teaching an introductory media course that covered various models of media persuasion, including the *opinion-leader* model. In this view, influence flows from the media to opinion leaders and then from opinion leaders to other people. Fortunately, there was an example on hand that many of my students knew about: the campaign to have Joseph Kony arrested for war crimes in Africa, including his use of child soldiers. This campaign was led by the organization Invisible Children, which targeted its message to young people in the West. In 2012, the group launched a social media campaign to have Kony

captured by the year's end. In addition to creating the online film *Kony 2012*, campaigners tweet-bombed celebrities; Rihanna, Justin Bieber, and Oprah retweeted for the cause. As media scholar and Internet activist Ethan Zuckerman notes, however, this "attention philanthropy" does not lend itself to indefinite repetition: "Oprah has a great deal of a valuable commodity—attention—and the incremental cost of her spending that attention to call attention to a cause is minimal. ... In the long run, if she tweets about every campaign her fans want her to promote, she'll likely start to lose her audience—the incremental cost may be small, but the cumulative cost could be very high."²⁴

A month after the Kony campaign, Bachir ("Athene") Boumaaza (an Internet celebrity known for his gaming prowess, YouTube channel, and social activism) attempted the same tactic as part of Operation ShareCraft. Athene and his collaborator Reese Leysen were attempting to leverage Athene's significant online following to ameliorate hunger in the Horn of Africa. Much like the Invisible Children campaign, Operation ShareCraft sought the attention of prominent social media personalities, including Xeni Jardin, a founding contributor of the popular blog *Boing Boing*. However, Jardin's attention was elsewhere. Four months earlier, she had scheduled an appointment for a mammogram and had decided that sharing her experience would allay her anxiety and perhaps encourage others:

I would tweet this new thing, like I do with lots of new things, and make the unknown and new feel less so. Maybe by doing so ... other women like me who'd never done this would also feel like it was less weird, less scary, more normal and worth doing without hesitation. I'd crack some 140-character jokes. I'd make fun of myself and others. I would Instagram my mammogram.²⁵

Apparently, this is not unusual. In his book about "social communication in the Twitter Age," Dhiraj Murthy dedicates a chapter to health, focusing on cancer. He writes that although many cancer-related tweets are either about charities or news of treatments, ordinary people share their own experiences, often at the intersection of the banal and the profound. People who were anxiously sitting in a doctor's waiting room were unlikely to write a blog entry about it but would "pick up their smartphone for 45 seconds and Tweet about it."²⁶

After Xení Jardin received the results of her mammogram, she tweeted “I have breast cancer. I am in good hands. There is a long road ahead and it leads to happiness and a cancer-free, long, healthy life.”²⁷ Many people found the reports of her ongoing treatment compelling, and fellow patients and survivors used Twitter to exchange information and support. Although Operation ShareCraft’s tweet-bomb plan was to ask celebrities “to spread the word or support the cause in any way possible,” Jardin found the group’s mass entreaties to be unseemly.²⁸ She tweeted to Athene and Reese that this was “totally tone-deaf and inappropriate” and complained of “Getting tons of SRY U HAZ CANSUR PLZ DON’T DIE XENI PS HELP US END HUNGER IN TEH HORN OF AFRICA! KTHXBYE spam tweets. Fuck all of you, srsly.”²⁹ A slew of hateful and misogynist tweets followed. Athene and others apologized, explaining that “trolls tried to sabotage the event,” yet Jardin continued to receive messages to the effect that “All these people sacrifice their free time to raise awareness and do good. If you dont wanna get tweet bombed dont be on Twitter.”³⁰ Jardin concluded that she had no ill will toward the charity and that this had been the “Strangest griefer/troll storm I’ve ever experienced here. Hope it’s not indicator of Twitter’s future.”

Comment in Context

This book isn’t about the future of Twitter, blogs, or YouTube, but it is about comment in the age of the Web. (Nor is this book about looking for comments from time-travelers from the future, as some researchers have done!³¹) Similarly, it’s not about how social media have transformed politics, journalism, or global relations. This book is about the stuff in the margins—the things that ordinary people encounter in daily life. YouTube’s Jawed Karim recognized the importance of user-generated content at HOTorNOT, but many continue to overlook the comments on such content. Indeed, the unsavory aspects of online comment have prompted many to turn a blind eye to the “bottom half of the Internet” and to advise visitors, “Don’t read the comments.” And although headlines publicize the IPOs and purchases of related Websites, many are unaware of the illicit markets in which followers, reviewers, and commenters are bought and sold.

In addition to defining comment as reactive, short, and asynchronous, it also is useful to consider its context. A comment is about something—an *object* or a topic, such as a book. A comment has a *source* or author, who might be

identified or anonymous. It has an *audience*—the people who are the intended readers of the comment. The *content* of a comment might be prose, a verbal aside, or a rating. Even clicking a *+1* or a *like* button is a comment. Finally, the *intentions* and *effects* of comments are important. A comment can affect another's status (for example, a diploma is a comment about academic standing, and gossip is a comment about social standing), it can help others make decisions (such as "The food here is great"), or it can alter a person's behavior (for example, by providing feedback about someone's actions). Much drama can ensue when the context of comments is ambiguous or transgressed, such as when a note to a friend at school is intercepted and becomes known to everyone in the classroom.

Two chapters in this book focus on the intentions and effects of comment: informing (such as with a review that helps readers choose a product or service) and providing feedback (such as with remarks that help people improve their lives). With respect to reviews, in chapter 2, I note that comment today inherits many reviewing modes from the past. The Web contains the legacies of early twentieth-century engineers who were preoccupied with comparative analysis and of two brothers who sought to sell more tires by assigning a number of stars to hotels and restaurants. I also discuss the *likers*, who share recommendations that are rooted in love and experience; the crowd, which shares its particular, peculiar, kind of wisdom; and the critics, who highlight and connect with analysis and insight. All forms of writing that have gone before are present on the Web—and at a very large scale. These types of comment existed before the twenty-first century, but never were they available in such great numbers or were they as easily accessible as they are today.

Because comments have become one of the most valued commodities of the day, they are subject to much manipulation. In chapter 3, I review the research on fake reviews and discuss some prominent cases of fakery. I distinguish between makers, fakers, and takers and discuss the dynamics and techniques of online fakery. I argue that relying on social networks (consuming the recommendations of our friends rather than strangers) will not solve the problem of manipulation but will tempt us to become manipulators ourselves.

In chapter 4, I focus on feedback, which is a practice that has great potential to go awry. This is especially so in the case of what I call *tweak critique*, such as altering someone else's photo to improve its composition. Because it is now easy to solicit and give feedback, quick and unthinking comment can easily bruise and can blur the distinctions between criticism, feedback, and review,

which result in contention and controversy. Similarly, the scope and scale of feedback have changed: feedback to one person can be seen by many and can be unsolicited and unwanted. Feedback requires the giver to be careful and the receiver to perform what scholars call "emotion work."

Unfortunately, online conflicts are not always between well-intended commenters who try to be civil. Although comment is a type of communication that permits us to be helpful, friendly, and encouraging to others, it can lead to feelings of frustration and alienation. In chapter 5, I describe the trolls and haters, bully battles, and misogyny that often are encountered online and frame this culture with what is known about the effects of anonymity, deindividuation, and depersonalization. Alienating and hateful comment is not likely to go away anytime soon, but we are beginning to give greater attention to how to respond to it at technical and communal levels.

Online comment is reactive and short, and these characteristics affect writers and readers in a couple of ways. First, our reaction to things (be it a comment, an answer to a question, or the liking of a photo) has come to be seen as a way in which we define ourselves. And the ways that others react to those reactions (such as by retweeting them) are seen as a valuation of those selves. Second, comment's shortness and ubiquity mean that attention is easily and often drawn online. And the fact that these comments can be counted and tracked also affects how people value themselves and others. Chapter 6 asks how this is shaping us. How does the nonstop stream of our own and others' photographs and status updates affect self-esteem and well-being? Are the short and asynchronous bursts of comment that are processed throughout the day affecting the ability to concentrate? Have the quantification and ranking of social relations gone too far? These are complex questions without easy answers.

Before concluding the book, I show in chapter 7 how puzzling comment is. In addition to informing, improving, alienating, manipulating, and shaping, these short and asynchronous messages can bemuse: they can be slapdash, confusing, amusing, revealing, and weird. Because comment is reactive, it is inherently contextual. Yet it also is *hypotextual* (that is, undertextual), shedding context with ease and prompting the comment of *WTF?!?* (what the fuck?). From this confusion and weirdness, however, we can learn about the advantages of moving first, the intricacies of communication, the science of rating systems, and the challenges of not losing context at the bottom of the Web.

[[Figure 1.1. Comp/AU: Please note that fig nos. are for ID only. They will not print.]]

© 2009 geek-and-poke.com CC-BY-3.0

[[Figure 1.2]]

© 2011 geek-and-poke.com CC-BY-3.0

{notes begin}

Chapter 1: Comment

1. John Cloud, “The YouTube Gurus,” *Time*, December 25, 2006, <http://www.time.com/time/printout/0,8816,1570795,00.html>.
2. xxxCRAZYchannel, “Darth Vader Breathing Ten Hours,” *YouTube*, June 17, 2012, <http://www.youtube.com/watch?v=un8FAjXWOBY>.
3. AvoidComments, “Don’t Read Comments,” *Twitter*, November 28, 2012, <https://twitter.com/AvoidComments>.
4. Marshal McLuhan, *Counterblast* (Toronto: McClelland and Steward, 1969), 5.
5. “Facebook’s Growth in the Past Year,” *Facebook*, 2013, <https://www.facebook.com/photo.php?fbid=10151908376831729>; Doug Gross, “On Facebook, Click ‘Like’ Can Help Scammers,” *CNN*, January 20, 2014, <http://www.cnn.com/2014/01/21/tech/social-media/facebook-like-farming/index.html>.
6. Robin Dunbar, *Grooming, Gossip, and the Evolution of Language* (Cambridge: Harvard University Press, 1996); Robin Dunbar, “Gossip in Evolutionary Perspective,” *Review of General Psychology* 8, no. 2 (2004): 100–110.

7. Alan Dundes, "Here I Sit: A Study of American Latrinalia," *Kroeber Anthropological Society Papers* 34 (1966): 91–105; Adam Trahan, "Identity and Ideology: The Dialogic Nature of Latrinalia," *Internet Journal of Criminology* (2011): 1–9; Robert F. Goodman and Aaron Ben-Ze'ev, eds., *Good Gossip* (Lawrence: University Press of Kansas, 1994); Aaron Ben-Ze'ev, "The Vindication of Gossip," in *Good Gossip*, ed. Robert F. Goodman and Aaron Ben-Ze'ev, 11–24 (Lawrence: University Press of Kansas, 1994); Nicholas DiFonzo and Prashant Bordia, "Rumor, Gossip and Urban Legends," *Diogenes* 54, no. 19 (2007): 28, doi:10.1177/0392192107073433.
8. Dunbar, *Grooming, Gossip, and the Evolution of Language*, 71.
9. Michael Erard, "No Comments," *New York Times Magazine*, September 20, 2013, <http://www.nytimes.com/2013/09/22/magazine/no-comments.html>.
10. John Markoff, "An Internet Critic Who Is Not Shy about Ruffling the Big Names in High Technology," *New York Times*, April 9, 2001, <http://www.nytimes.com/2001/04/09/business/internet-critic-who-not-shy-about-ruffling-big-names-high-technology.html>.
11. Dave Winer, "The Unedited Voice of a Person," *Scripting News*, January 1, 2007, <http://scripting.com/2007/01/01.html#theUneditedVoiceOfAPerson>.
12. Dave Winer, "Proposal: A New Kind of Blog Comment System," *Scripting News*, August 22, 2010, <http://scripting.com/stories/2010/08/22/proposalANewKindOfBlogComm.html>; Dave Winer, "No Comment," *Scripting News*, February 19, 2012, <http://scripting.com/stories/2012/02/19/noComment.html>.
13. Mathew Wingram, "Facebook: Dave Winer Brings the Hate," *Work Blog*, October 14, 2007, <http://www.mathewwingram.com/work/2007/10/14/facebook-dave-winer-brings-the-hate/>.
14. Newshour, "Post Web Site Halts Comments Section," *PBS*, January 24, 2006, http://www.pbs.org/newshour/bb/media/jan-june06/post_1-24.html; Joshua Topolsky, "We're Turning Comments Off for a Bit," *Engadget*, February 2, 2010, <http://www.engadget.com/2010/02/02/were-turning-comments-off-for-a-bit/>;

Suzanne Labarre, “Why We’re Shutting Off Our Comments,” *Popular Science*, September 24, 2013, <http://www.popsci.com/science/article/2013-09/why-were-shutting-our-comments?src=SOC&dom=tw>; Rob Beschizza, “Time to Turn Off Comments for Good,” *Muckrack*, February 28, 2013, <http://muckrack.com/Beschizza/statuses/307172805252894722>.

15. “Introducing Our Comments Box Launch Partners,” *Facebook*, March 3, 2011, http://www.facebook.com/note.php?note_id=187369061298591&comments; Farhad Manjoo, “Anonymous Comments: Why We Need to Get Rid of Them Once and for All,” *Anonymous Comments*, March 9, 2011, http://www.slate.com/articles/technology/technology/2011/03/troll_reveal_thyself.html.

16. MG Siegler, “Facebook Comments Have Silenced the Trolls—But Is It Too Quiet?,” *TechCrunch*, March 6, 2011, <http://techcrunch.com/2011/03/06/techcrunch-facebook-comments/>.

17. Clay Shirky, “Group as User: Flaming and the Design of Social Software,” *Shirky*, November 5, 2004, http://www.shirky.com/writings/group_user.html; Clay Shirky, “‘Hottubbing’ as a Community Management Pattern,” *Many-to-Many*, May 15, 2003, http://many.corante.com/archives/2003/05/15/hottubbing_as_a_community_management_pattern.php; Clay Shirky, *Cognitive Surplus: Creativity and Generosity in a Connected Age* (New York: Penguin Press, 2010), 197–202; Clay Shirky, “Why Do Comments Suck?,” *Microsoft Research*, January 17, 2013, <http://research.microsoft.com/apps/video/dl.aspx?id=180323>.

18. Lawrence Lessig, “Announcing the Hibernation of Lessig.org/blog,” *Blog*, August 20, 2009, <http://www.lessig.org/2009/08/>.

19. Alice Marwick and danah boyd, “To See and Be Seen: Celebrity Practice on Twitter,” *Convergence* 17, no. 2 (May 2011): 139–158.

20. Mor Naaman, Jeffrey Boase, and Chih-Hui Lai, “Is It Really about Me?,” in *Proceedings of CSCW 2010*, 2010, <http://infolab.stanford.edu/~mor/research/naamanCSCW10.pdf>.

21. Diana I. Tamir and Jason P. Mitchell, “Disclosing Information about the Self Is Intrinsically Rewarding,” *Proceedings of the National Academy of Sciences* 109 (2012): 8038–8043, doi:10.1073/pnas.1202129109.
22. Trent Reznor, “Online Communities, Etc.,” *NIN Forums*, June 10, 2009, <http://forum.nin.com/bb/read.php?59,731489>.
23. Admin, “The Atheism+ Block Bot,” *The Blockbot*, March 7, 2013, <http://www.theblockbot.com/>.
24. Ethan Zuckerman, “The Tweetbomb and the Ethics of Attention,” *My Heart’s in Accra*, April 20, 2012, <http://www.ethanzuckerman.com/blog/2012/04/20/the-tweetbomb-and-the-ethics-of-attention>.
25. Xeni Jardin, “The Diagnosis,” *Boing Boing*, December 9, 2011, <http://boingboing.net/2011/12/09/the-diagnosis.html>.
26. Dhiraj Murthy, *Twitter: Social Communication in the Twitter Age* (Malden, MA: Polity, 2013), 140.
27. Xeni Jardin, “Twitter / Xeni: I Have Breast Cancer. I Am ...,” *Twitter*, December 1, 2011, <https://twitter.com/xeni/status/142437402626105344>.
28. Reese Leysen, “ShareCraft Celebrity Tweet-Bombing (and Other Ideas),” *I Power*, April 6, 2012, <http://ipowerproject.com/forum/topics/sharecraft-celebrity-tweet-bombing-and-other-ideas>.
29. Zuckerman, “The Tweetbomb and the Ethics of Attention.”
30. mc rider, “@xeni All These People Sacrifice ...,” *Twitter*, April 19, 2012, <https://twitter.com/wesir23/status/193089878580400128>.
31. Robert J. Nemiroff and Teresa Wilson, “Searching the Internet for Evidence of Time Travelers,” *Arxiv*, December 26, 2013, <http://arxiv.org/abs/1312.7128>.

{notes end}