

make sense of his life. In fact, when Winston thought about it, he realized there were a lot of other e-mails from his life that fit into this odd category—stuff you don't want to look at but don't want to lose, either. So he took all these emotionally difficult messages and archived them in Gmail using an evocative label: "Forget." Out of sight, out of mind, but retrievable.

It's a beautiful metaphor for the odd paradoxes and trade-offs we'll live with in a world of infinite memory. Our ancestors learned how to remember; we'll learn how to forget.

Public Thinking

In 2003, Kenyan-born Ory Okolloh was a young law student who was studying in the United States but still obsessed with Kenyan politics. There was plenty to obsess over. Kenya was a cesspool of government corruption, ranking near the dismal bottom on the Corruption Perceptions Index. Okolloh spent hours and hours talking to her colleagues about it, until eventually one suggested the obvious: *Why don't you start a blog?*

Outside of essays for class, she'd never written anything for an audience. But she was game, so she set up a blog and faced the keyboard.

"I had zero ideas about what to say," she recalls.

This turned out to be wrong. Over the next seven years, Okolloh revealed a witty, passionate voice, keyed perfectly to online conversation. She wrote a steady stream of posts about the battle against Kenyan corruption, linking to reports of bureaucrats spending enormous sums on luxury vehicles and analyzing the "Anglo-leasing scandal," in which the government paid hundreds of millions for services—like producing a new passport system for the country—that were never delivered. When she moved back to Kenya in 2006, she began posting snapshots of such things as the bathtub-sized muddy potholes on the road to the airport. ("And our economy is supposed to be growing how exactly?") Okolloh also wrote about daily life, posting pictures of her baby and discussing the joys of living in Nairobi, including cabdrivers so friendly they'd run errands

for her. She gloated nakedly when the Pittsburgh Steelers, her favorite football team, won a game.

After a few years, she'd built a devoted readership, including many Kenyans living in and out of the country. In the comments, they'd joke about childhood memories like the "packed lunch trauma" of low-income kids being sent to school with ghastly leftovers. Then in 2007, the ruling party rigged the national election and the country exploded in violence. Okolloh wrote anguished posts, incorporating as much hard information as she could get. The president imposed a media blackout, so the country's patchy Internet service was now a crucial route for news. Her blog quickly became a clearinghouse for information on the crisis, as Okolloh posted into the evening hours after coming home from work.

"I became very disciplined," she tells me. "Knowing I had these people reading me, I was very self-conscious to build my arguments, back up what I wanted to say. It was very interesting; I got this sense of obligation."

Publishers took notice of her work and approached Okolloh to write a book about her life. She turned them down. The idea terrified her. A whole book? "I have a very introverted real personality," she adds.

Then one day a documentary team showed up to interview Okolloh for a film they were producing about female bloggers. They'd printed up all her blog posts on paper. When they handed her the stack of posts, it was the size of two telephone books.

"It was huge! Humongous!" She laughs. "And I was like, *oh my*. That was the first time I had a sense of the volume of it." Okolloh didn't want to write a book, but in a sense, she already had.



The Internet has produced a foaming Niagara of writing. Consider these current rough estimates: Each day, we compose 154 billion

e-mails, more than 500 million tweets on Twitter, and over 1 million blog posts and 1.3 million blog comments on WordPress alone. On Facebook, we write about 16 billion words per day. That's just in the United States: in China, it's 100 million updates each day on Sina Weibo, the country's most popular microblogging tool, and millions more on social networks in other languages worldwide, including Russia's VK. Text messages are terse, but globally they're our most frequent piece of writing: 12 billion per day.

How much writing is that, precisely? Well, doing an extraordinarily crude back-of-the-napkin calculation, and sticking only to e-mail and utterances in social media, I calculate that we're composing at least 3.6 trillion words daily, or the equivalent of 36 million books every day. The entire U.S. Library of Congress, by comparison, holds around about 35 million books.

I'm not including dozens of other genres of online composition, each of which comprises entire subgalaxies of writing, because I've never been able to find a good estimate of their size. But the numbers are equally massive. There's the world of fan fiction, the subculture in which fans write stories based on their favorite TV shows, novels, manga comics, or just about anything with a good story world and cast of characters. When I recently visited Fanfiction.net, a large repository of such writing, I calculated—again, using some equally crude napkin estimates—that there were about 325 million words' worth of stories written about the popular young-adult novel *The Hunger Games*, with each story averaging around fourteen thousand words. That's just for one book: there are thousands of other forums crammed full of writing, ranging from twenty-six thousand *Star Wars* stories to more than seventeen hundred pieces riffing off Shakespeare's works. And on top of fan fiction, there are also all the discussion boards, talmudically winding comment threads on blogs and newspapers, sprawling wikis, meticulously reported recaps of TV shows, or blow-by-blow walk-through dissections of video

games; some of the ones I've used weigh in at around forty thousand words. I would hazard we're into the trillions now.

Is any of this writing good? Well, that depends on your standards, of course. I personally enjoyed Okoloh's blog and am regularly astonished by the quality and length of expression I find online, the majority of which is done by amateurs in their spare time. But certainly, measured against the prose of an Austen, Orwell, or Tolstoy, the majority of online publishing pales. This isn't surprising. The science fiction writer Theodore Sturgeon famously said something like, "Ninety percent of everything is crap," a formulation that geeks now refer to as Sturgeon's Law. Anyone who's spent time slogging through the swamp of books, journalism, TV, and movies knows that Sturgeon's Law holds pretty well even for edited and curated culture. So a global eruption of unedited, everyday self-expression is probably even more likely to produce this 90-10 split—an ocean of dreck, dotted sporadically by islands of genius. Nor is the volume of production uniform. Surveys of commenting and posting generally find that a minority of people are doing most of the creation we see online. They're ferociously overproductive, while the rest of the online crowd is quieter. Still, even given those parameters and limitations, the sheer profusion of thoughtful material that is produced every day online is enormous.

And what makes this explosion truly remarkable is what came before: comparatively little. For many people, almost nothing.

Before the Internet came along, most people rarely wrote anything at all for pleasure or intellectual satisfaction after graduating from high school or college. This is something that's particularly hard to grasp for professionals whose jobs require incessant writing, like academics, journalists, lawyers, or marketers. For them, the act of writing and hashing out your ideas seems commonplace. But until the late 1990s, this simply wasn't true of the average nonliterary person. The one exception was the white-collar workplace, where

jobs in the twentieth century increasingly required more memo and report writing. But personal expression outside the workplace—in the curious genres and epic volume we now see routinely online—was exceedingly rare. For the average person there were few vehicles for publication.

What about the glorious age of letter writing? The reality doesn't match our fond nostalgia for it. Research suggests that even in the United Kingdom's peak letter-writing years—the late nineteenth century, before the telephone became common—the average citizen received barely one letter every two weeks, and that's even if we generously include a lot of distinctly unliterary business missives of the "hey, you owe us money" type. (Even the ultraliterate elites weren't pouring out epistles. They received on average two letters per week.) In the United States, the writing of letters greatly expanded after 1845, when the postal service began slashing its rates on personal letters and an increasingly mobile population needed to communicate across distances. Cheap mail was a powerful new mode of expression—though as with online writing, it was unevenly distributed, with probably only a minority of the public taking part fully, including some city dwellers who'd write and receive mail every day. But taken in aggregate, the amount of writing was remarkably small by today's standards. As the historian David Henkin notes in *The Postal Age*, the per capita volume of letters in the United States in 1860 was only 5.15 per year. "That was a huge change at the time—it was important," Henkin tells me. "But today it's the exceptional person who doesn't write five messages a day. I think a hundred years from now scholars will be swimming in a bewildering excess of life writing."

As an example of the pre-Internet age, consider my mother. She's seventy-seven years old and extremely well read—she received a terrific education in the Canadian high school system and voraciously reads novels and magazines. But she doesn't use the Internet to

express herself; she doesn't write e-mail, comment on discussion threads or Facebook, post status updates, or answer questions online. So I asked her how often in the last year she'd written something of at least a paragraph in length. She laughed. "Oh, never!" she said. "I sign my name on checks or make lists—that's about it." Well, how about in the last ten years? Nothing to speak of, she recalled. I got desperate: How about twenty or thirty years back? Surely you wrote letters to family members? Sure, she said. But only about "three or four a year." In her job at a rehabilitation hospital, she jotted down the occasional short note about a patient. You could probably take all the prose she's generated since she left high school in 1952 and fit it in a single file folder.

Literacy in North America has historically been focused on reading, not writing; consumption, not production. Deborah Brandt, a scholar who researched American literacy in the 1980s and '90s, has pointed out a curious aspect of parenting: while many parents worked hard to ensure their children were regular readers, they rarely pushed them to become regular writers. You can understand the parents' point of view. In the industrial age, if you happened to write something, you were extremely unlikely to publish it. Reading, on the other hand, was a daily act crucial for navigating the world. Reading is also understood to have a moral dimension; it's supposed to make you a better person. In contrast, Brandt notes, writing was something you did mostly for work, serving an industrial purpose and not personal passions. Certainly, the people Brandt studied often enjoyed their work writing and took pride in doing it well. But without the impetus of the job, they wouldn't be doing it at all. Outside of the office, there were fewer reasons or occasions to do so.

The advent of digital communications, Brandt argues, has upended that notion. We are now a global culture of avid writers. Some of this boom has been at the workplace; the clogged e-mail

inboxes of white-collar workers testifies to how much for-profit verbiage we crank out. But in our own time, we're also writing a stunning amount of material about things we're simply interested in—our hobbies, our friends, weird things we've read or seen online, sports, current events, last night's episode of our favorite TV show. As Brandt notes, reading and writing have become blended: "People read in order to generate writing; we read from the posture of the writer; we write to other people who write." Or as Francesca Coppa, a professor who studies the enormous fan fiction community, explains to me, "It's like the Bloomsbury Group in the early twentieth century, where everybody is a writer and everybody is an audience. They were all writers who were reading each other's stuff, and then writing about that, too."

We know that reading changes the way we think. Among other things, it helps us formulate thoughts that are more abstract, categorical, and logical.

So how is all this *writing* changing our cognitive behavior?

For one, it can help clarify our thinking.

Professional writers have long described the way that the act of writing forces them to distill their vague notions into clear ideas. By putting half-formed thoughts on the page, we externalize them and are able to evaluate them much more objectively. This is why writers often find that it's only when they start writing that they figure out what they want to say.

Poets famously report this sensation. "I do not sit down at my desk to put into verse something that is already clear in my mind," Cecil Day-Lewis wrote of his poetic compositions. "If it were clear in my mind, I should have no incentive or need to write about it. . . . We do not write in order to be understood; we write in order to understand." William Butler Yeats originally intended "Leda and

the Swan” to be an explicitly political poem about the impact of Hobbesian individualism; in fact, it was commissioned by the editor of a political magazine. But as Yeats played around on the page, he became obsessed with the existential dimensions of the Greek myth of Leda—and the poem transformed into a spellbinding meditation on the terrifying feeling of being swept along in forces beyond your control. “As I wrote,” Yeats later recalled, “bird and lady took such possession of the scene that all politics went out of it.” This phenomenon isn’t limited to poetry. Even the workplace that Brandt studied—including all those memos cranked out at white-collar jobs—help clarify one’s thinking, as many of Brandt’s subjects told her. “It crystallizes you,” one said. “It crystallizes your thought.”

The explosion of online writing has a second aspect that is even more important than the first, though: it’s almost always done for an *audience*. When you write something online—whether it’s a one-sentence status update, a comment on someone’s photo, or a thousand-word post—you’re doing it with the expectation that someone might read it, even if you’re doing it anonymously.

Audiences clarify the mind even more. Bloggers frequently tell me that they’ll get an idea for a blog post and sit down at the keyboard in a state of excitement, ready to pour their words forth. But pretty soon they think about the fact that someone’s going to *read* this as soon as it’s posted. And suddenly all the weak points in their argument, their clichés and lazy, autofill thinking, become painfully obvious. Gabriel Weinberg, the founder of DuckDuckGo—an upstart search engine devoted to protecting its users’ privacy—writes about search-engine politics, and he once described the process neatly:

Blogging forces you to write down your arguments and assumptions. This is the single biggest reason to do it, and I think it alone makes it worth it. You have a lot of opinions. I’m sure some of them you hold strongly. Pick

one and write it up in a post—I’m sure your opinion will change somewhat, or at least become more nuanced. When you move from your head to “paper,” a lot of the hand-waveyness goes away and you are left to really defend your position to yourself.

“Hand waving” is a lovely bit of geek coinage. It stands for the moment when you try to show off to someone else a cool new gadget or piece of software you created, which suddenly won’t work. Maybe you weren’t careful enough in your wiring; maybe you didn’t calibrate some sensor correctly. Either way, your invention sits there broken and useless, and the audience stands there staring. In a panic, you try to *describe* how the gadget works, and you start waving your hands to illustrate it: hand waving. But nobody’s ever convinced. Hand waving means you’ve failed. At MIT’s Media Lab, the students are required to show off their new projects on Demo Day, with an audience of interested spectators and corporate sponsors. For years the unofficial credo was “demo or die”: if your project didn’t work as intended, you died (much as stand-up comedians “die” on stage when their act bombs). I’ve attended a few of these events and watched as some poor student’s telepresence robot freezes up and crashes . . . and the student’s desperate, white-faced hand waving begins.

When you walk around meditating on an idea quietly to yourself, you do a lot of hand waving. It’s easy to win an argument inside your head. But when you face a *real* audience, as Weinberg points out, the hand waving has to end. One evening last spring he rented the movie *Moneyball*, watching it with his wife after his two toddlers were in bed. He’s a programmer, so the movie—about how a renegade baseball coach picked powerful players by carefully analyzing their statistics—inspired five or six ideas he wanted to blog about the next day. But as usual, those ideas were rather fuzzy, and

it wasn't until he sat down at the keyboard that he realized he wasn't quite sure what he was trying to say. He was hand waving.

"Even if I was publishing it to no one, it's just the *threat* of an audience," Weinberg tells me. "If someone could come across it under my name, I have to take it more seriously." Crucially, he didn't want to bore anyone. Indeed, one of the unspoken cardinal rules of online expression is *be more interesting*—the sort of social pressure toward wit and engagement that propelled coffeehouse conversations in Europe in the nineteenth century. As he pecked away at the keyboard, trying out different ideas, Weinberg slowly realized what interested him most about the movie. It wasn't any particularly clever bit of math the baseball coach had performed. No, it was how the coach's focus on numbers had created a new way to excel at baseball. The baseball coach's behavior reminded him of how small entrepreneurs succeed: they figure out something that huge, intergalactic companies simply can't spot, because they're stuck in their old mind-set. Weinberg's process of crafting his idea—and trying to make it clever for his readers—had uncovered its true dimensions. Reenergized, he dashed off the blog entry in a half hour.

Social scientists call this the "audience effect"—the shift in our performance when we know people are watching. It isn't always positive. In live, face-to-face situations, like sports or live music, the audience effect often makes runners or musicians perform better, but it can sometimes psych them out and make them choke, too. Even among writers I know, there's a heated divide over whether thinking about your audience is fatal to creativity. (Some of this comes down to temperament and genre, obviously: Oscar Wilde was a brilliant writer and thinker who spent his life swanning about in society, drawing the energy and making the observations that made his plays and essays crackle with life; Emily Dickinson was a

brilliant writer and thinker who spent her life sitting at home alone, quivering neurasthenically.)

But studies have found that particularly when it comes to analytic or critical thought, the effort of communicating to someone else forces you to think more precisely, make deeper connections, and learn more.

You can see this audience effect even in small children. In one of my favorite experiments, a group of Vanderbilt University professors in 2008 published a study in which several dozen four- and five-year-olds were shown patterns of colored bugs and asked to predict which would be next in the sequence. In one group, the children simply solved the puzzles quietly by themselves. In a second group, they were asked to explain into a tape recorder how they were solving each puzzle, a recording they could keep for themselves. And in the third group, the kids had an audience: they had to explain their reasoning to their mothers, who sat near them, listening but not offering any help. Then each group was given patterns that were more complicated and harder to predict.

The results? The children who solved the puzzles silently did worst of all. The ones who talked into a tape recorder did better—the mere act of articulating their thinking process aloud helped them think more critically and identify the patterns more clearly. But the ones who were talking to a meaningful audience—Mom—did best of all. When presented with the more complicated puzzles, on average they solved more than the kids who'd talked to themselves and about twice as many as the ones who'd worked silently.

Researchers have found similar effects with older students and adults. When asked to write for a real audience of students in another country, students write essays that are substantially longer and have better organization and content than when they're writing for their teacher. When asked to contribute to a wiki—a space that's

highly public and where the audience can respond by deleting or changing your words—college students snap to attention, writing more formally and including more sources to back up their work. Brenna Clarke Gray, a professor at Douglas College in British Columbia, assigned her English students to create Wikipedia entries on Canadian writers, to see if it would get them to take the assignment more seriously. She was stunned how well it worked. “Often they’re handing in these short essays without any citations, but with Wikipedia they suddenly were staying up to two a.m. honing and rewriting the entries and carefully sourcing everything,” she tells me. The reason, the students explained to her, was that their audience—the Wikipedia community—was quite gimlet eyed and critical. They were harder “graders” than Gray herself. When the students first tried inputting badly sourced articles, the Wikipedians simply deleted them. So the students were forced to go back, work harder, find better evidence, and write more persuasively. “It was like night and day,” Gray adds.

Sir Francis Bacon figured this out four centuries ago, quipping that “reading maketh a full man, conference a ready man, and writing an exact man.”

Interestingly, the audience effect doesn’t necessarily require a big audience to kick in. This is particularly true online. Weinberg, the DuckDuckGo blogger, has about two thousand people a day looking at his blog posts; a particularly lively response thread might only be a dozen comments long. It’s not a massive crowd, but from his perspective it’s transformative. In fact, many people have told me they feel the audience effect kick in with even a tiny handful of viewers. I’d argue that the cognitive shift in going from an audience of zero (talking to yourself) to an audience of ten people (a few friends or random strangers checking out your online post) is so big that it’s actually huger than going from ten people to a million people.

This is something that the traditional thinkers of the industrial

age—particularly print and broadcast journalists—have trouble grasping. For them, an audience doesn’t mean anything unless it’s massive. If you’re writing specifically to make money, you need a large audience. An audience of ten is meaningless. Economically, it means you’ve failed. This is part of the thinking that causes traditional media executives to scoff at the spectacle of the “guy sitting in his living room in his pajamas writing what he thinks.” But for the rest of the people in the world, who never did much nonwork writing in the first place—and who almost never did it for an audience—even a handful of readers can have a vertiginous, catalytic impact.

Writing about things has other salutary cognitive effects. For one, it improves your memory: write about something and you’ll remember it better, in what’s known as the “generation effect.” Early evidence came in 1978, when two psychologists tested people to see how well they remembered words that they’d written down compared to words they’d merely read. Writing won out. The people who wrote words remembered them better than those who’d only read them—probably because generating text yourself “requires more cognitive effort than does reading, and effort increases memorability,” as the researchers wrote. College students have harnessed this effect for decades as a study technique: if you force yourself to jot down what you know, you’re better able to retain the material.

This sudden emergence of audiences is significant enough in Western countries, where liberal democracies guarantee the right to free speech. But in countries where there’s less of a tradition of free speech, the emergence of networked audiences may have an even more head-snapping effect. When I first visited China to meet some of the country’s young bloggers, I’d naively expected that most of them would talk about the giddy potential of arguing about human rights and free speech online. I’d figured that for people living in an authoritarian country, the first order of business, once you had a public microphone, would be to agitate for democracy.

But many of them told me it was startling enough just to suddenly be writing, in public, about the minutiae of their everyday lives—arguing with friends (and interested strangers) about stuff like whether the movie *Titanic* was too sappy, whether the fashion in the *Super Girl* competitions was too racy, or how they were going to find jobs. “To be able to speak about what’s going on, what we’re watching on TV, what books we’re reading, what we feel about things, that is a remarkable feeling,” said a young woman who had become Internet famous for writing about her sex life. “It is completely different from what our parents experienced.” These young people believed in political reform, too. But they suspected that the creation of small, everyday audiences among the emerging middle-class online community, for all the seeming triviality of its conversation, was a key part of the reform process.

Once thinking is public, connections take over. Anyone who’s googled their favorite hobby, food, or political subject has immediately discovered that there’s some teeming site devoted to servicing the infinitesimal fraction of the public that shares their otherwise wildly obscure obsession. (Mine: building guitar pedals, modular origami, and the 1970s anime show *Battle of the Planets*). Propelled by the hyperlink—the ability of anyone to link to anyone else—the Internet is a connection-making machine.

And making connections is a big deal in the history of thought—and its future. That’s because of a curious fact: If you look at the world’s biggest breakthrough ideas, they often occur simultaneously to different people.

This is known as the theory of multiples, and it was famously documented in 1922 by the sociologists William Ogburn and Dorothy Thomas. When they surveyed the history of major modern in-

ventions and scientific discoveries, they found that almost all the big ones had been hit upon by different people, usually within a few years of each other and sometimes within a few weeks. They cataloged 148 examples: Oxygen was discovered in 1774 by Joseph Priestley in London and Carl Wilhelm Scheele in Sweden (and Scheele had hit on the idea several years earlier). In 1610 and 1611, four different astronomers—including Galileo—independently discovered sunspots. John Napier and Henry Briggs developed logarithms in Britain while Joost Bürgi did it independently in Switzerland. The law of the conservation of energy was laid claim to by four separate people in 1847. And radio was invented at the same time around 1900 by Guglielmo Marconi and Nikola Tesla.

Why would the same ideas occur to different people at the same time? Ogburn and Thomas argued that it was because our ideas are, in a crucial way, partly products of our environment. They’re “inevitable.” When they’re ready to emerge, they do. This is because we, the folks coming up with the ideas, do not work in a sealed-off, Rodin’s *Thinker* fashion. The things we think about are deeply influenced by the state of the art around us: the conversations taking place among educated folk, the shared information, tools, and technologies at hand. If four astronomers discovered sunspots at the same time, it’s partly because the quality of lenses in telescopes in 1611 had matured to the point where it was finally possible to pick out small details on the sun and partly because the question of the sun’s role in the universe had become newly interesting in the wake of Copernicus’s heliocentric theory. If radio was developed at the same time by two people, that’s because the basic principles that underpin the technology were also becoming known to disparate thinkers. Inventors knew that electricity moved through wires, that electrical currents caused fields, and that these seemed to be able to jump distances through the air. With that base of knowledge,


curious minds are liable to start wondering: Could you use those signals to communicate? And as Ogburn and Thomas noted, there are a *lot* of curious minds. Even if you assume the occurrence of true genius is pretty low (they estimated that one person in one hundred was in the “upper tenth” for smarts), that’s still a heck of a lot of geniuses.

When you think of it that way, what’s strange is not that big ideas occurred to different people in different places. What’s strange is that this didn’t happen *all the time, constantly*.


But maybe it did—and the thinkers just weren’t yet in contact. Thirty-nine years after Ogburn and Thomas, sociologist Robert Merton took up the question of multiples. (He’s the one who actually coined the term.) Merton noted an interesting corollary, which is that when inventive people aren’t aware of what others are working on, the pace of innovation slows. One survey of mathematicians, for example, found that 31 percent complained that they had needlessly duplicated work that a colleague was doing—because they weren’t aware it was going on. Had they known of each other’s existence, they could have collaborated and accomplished their calculations more quickly or with greater insight.

As an example, there’s the tragic story of Ernest Duchesne, the original discoverer of penicillin. As legend has it, Duchesne was a student in France’s military medical school in the mid-1890s when he noticed that the stable boys who tended the army’s horses did something peculiar: they stored their saddles in a damp, dark room so that mold would grow on their undersurfaces. They did this, they explained, because the mold helped heal the horses’ saddle sores. Duchesne was fascinated and conducted an experiment in which he treated sick guinea pigs with a solution made from mold—a rough form of what we’d now call penicillin. The guinea pigs healed completely. Duchesne wrote up his findings in a PhD thesis,

but because he was unknown and young—only twenty-three at the time—the French Institut Pasteur wouldn’t acknowledge it. His research vanished, and Duchesne died fifteen years later during his military service, reportedly of tuberculosis. It would take another thirty-two years for Scottish scientist Alexander Fleming to rediscover penicillin, independently and with no idea that Duchesne had already done it. Untold millions of people died in those three decades of diseases that could have been cured. Failed networks kill ideas.

When you can resolve multiples and connect people with similar obsessions, the opposite happens. People who are talking and writing and working on the same thing often find one another, trade ideas, and collaborate. Scientists have for centuries intuited the power of resolving multiples, and it’s part of the reason that in the seventeenth century they began publishing scientific journals and setting standards for citing the similar work of other scientists. Scientific journals and citation were a successful attempt to create a worldwide network, a mechanism for not just thinking in public but doing so in a connected way. As the story of Duchesne shows, it works pretty well, but not all the time. 

Today we have something that works in the same way, but for everyday people: the Internet, which encourages public thinking and resolves multiples on a much larger scale and at a pace more dementedly rapid. It’s now the world’s most powerful engine for putting heads together. Failed networks kill ideas, but successful ones trigger them.

As an example of this, consider what happened next to Ory Okol-
le.  During the upheaval after the rigged Kenyan election of 2007, she began tracking incidents of government violence. People called

and e-mailed her tips, and she posted as many as she could. She wished she had a tool to do this automatically—to let anyone post an incident to a shared map. So she wrote about that:

Google Earth supposedly shows in great detail where the damage is being done on the ground. It occurs to me that it will be useful to keep a record of this, if one is thinking long-term. For the reconciliation process to occur at the local level the truth of what happened will first have to come out. Guys looking to do something—any techies out there willing to do a mashup of where the violence and destruction is occurring using Google Maps?

One of the people who saw Okolloh's post was Erik Hersman, a friend and Web site developer who'd been raised in Kenya and lived in Nairobi. The instant Hersman read it, he realized he knew someone who could make the idea a reality. He called his friend David Kobia, a Kenyan programmer who was working in Birmingham, Alabama. Much like Okolloh, Kobia was interested in connecting Kenyans to talk about the country's crisis, and he had created a discussion site devoted to it. Alas, it had descended into political toxicity and calls for violence, so he'd shut it down, depressed by having created a vehicle for hate speech. He was driving out of town to visit some friends when he got a call from Hersman. Hersman explained Okolloh's idea—a map-based tool for reporting violence—and Kobia immediately knew how to make it happen. He and Hersman contacted Okolloh, Kobia began frantically coding with them, and within a few days they were done. The tool allowed anyone to pick a location on a Google Map of Kenya, note the time an incident occurred, and describe what happened. They called it Ushahidi—the Swahili word for "testimony."

Within days, Kenyans had input thousands of incidents of elec-

toral violence. Soon after, Ushahidi attracted two hundred thousand dollars in nonprofit funds and the trio began refining it to accept reports via everything from SMS to Twitter. Within a few years, Ushahidi had become an indispensable tool worldwide, with governments and nonprofits relying on it to help determine where to send assistance. After a massive earthquake hit Haiti in 2010, a Ushahidi map, set up within hours, cataloged twenty-five thousand text messages and more than four million tweets over the next month. It has become what Ethan Zuckerman, head of MIT's Center for Civic Media, calls "one of the most globally significant technology projects."

The birth of Ushahidi is a perfect example of the power of public thinking and multiples. Okolloh could have simply wandered around wishing such a tool existed. Kobia could have wandered around wishing he could use his skills to help Kenya. But because Okolloh was thinking out loud, and because she had an audience of like-minded people, serendipity happened.

The tricky part of public thinking is that it works best in situations where people aren't worried about "owning" ideas. The existence of multiples—the knowledge that people out there are puzzling over the same things you are—is enormously exciting if you're trying to solve a problem or come to an epiphany. But if you're trying to make *money*? Then multiples can be a real problem. Because in that case you're trying to stake a claim to ownership, to being the first to think of something. Learning that other people have the same idea can be anything from annoying to terrifying.

Scientists themselves are hardly immune. Because they want the fame of discovery, once they learn someone else is working on a similar problem, they're as liable to *compete* as to collaborate—and they'll bicker for decades over who gets credit. The story of penicillin illustrates this as well. Three decades after Duchesne made his discovery of penicillin, Alexander Fleming in 1928 stumbled on it

again, when some mold accidentally fell into a petri dish and killed off the bacteria within. But Fleming didn't seem to believe his discovery could be turned into a lifesaving medicine, so, remarkably, he never did any animal experiments and soon after dropped his research entirely. Ten years later, a pair of scientists in Britain—Ernest Chain and Howard Florey—read about Fleming's work, intuited that penicillin could be turned into a medicine, and quickly created an injectable drug that cured infected mice. After the duo published their work, Fleming panicked: someone else might get credit for his discovery! He hightailed it over to Chain and Florey's lab, greeting them with a wonderfully undercutting remark: "I have come to see what you've been doing with *my* old penicillin." The two teams eventually worked together, transforming penicillin into a mass-produced drug that saved countless lives in World War II. But for years, even after they all received a Nobel Prize, they jostled gently over who ought to get credit.

The business world is even more troubled by multiples. It's no wonder; if you're trying to make some money, it's hardly comforting to reflect on the fact that there are hundreds of others out there with precisely the same concept. Patents were designed to prevent someone else from blatantly infringing on your idea, but they also function as a response to another curious phenomenon: *unintentional* duplication. Handing a patent on an invention to one person creates artificial scarcity. It is a crude device, and patent offices have been horribly abused in recent years by "patent trolls"; they're people who get a patent for something (either by conceiving the idea themselves, or buying it) without any intention of actually producing the invention—it's purely so they can sue, or soak, people who go to market with the same concept. Patent trolls employ the concept of multiples in a perverted reverse, using the common nature of new ideas to hold all inventors hostage.

I've talked to entrepreneurs who tell me they'd like to talk openly online about what they're working on. They *want* to harness multiples. But they're worried that someone will take their idea and execute it more quickly than they can. "I know I'd get better feedback on my project if I wrote and tweeted about it," one once told me, "but I can't risk it." This isn't universally true; some start-up CEOs have begun trying to be more open, on the assumption that, as Bill Joy is famously reported quipping, "No matter who you are, most of the smartest people work for someone else." They know that talking about a problem makes it more likely you'll hook up with someone who has an answer.

But on balance, the commercial imperative to "own" an idea explains why public thinking has been a boon primarily for everyday people (or academics or nonprofits) pursuing their amateur passions. If you're worried about making a profit, multiples dilute your special position in the market; they're depressing. But if you're just trying to improve your thinking, multiples are exciting and catalytic. Everyday thinkers online are thrilled to discover someone else with the same idea as them.

We can see this in the history of "giving credit" in social media. Every time a new medium for public thinking has emerged, early users set about devising cordial, Emily Post-esque protocols. The first bloggers in the late 1990s duly linked back to the sources where they'd gotten their fodder. They did it so assiduously that the creators of blogging software quickly created an automatic "track-back" tool to help automate the process. The same thing happened on Twitter. Early users wanted to hold conversations, so they began using the @ reply to indicate they were replying to someone—and then to credit the original user when retweeting a link or pithy remark. Soon the hashtag came along—like #stupidestthingivedone today or #superbowl—to create floating, ad hoc conversations. All

these innovations proved so popular that Twitter made them a formal element of its service. We so value conversation and giving credit that we hack it into any system that comes along.

Stanford University English professor Andrea Lunsford is one of America's leading researchers into how young people write. If you're worried that college students today can't write as well as in the past, her work will ease your mind. For example, she tracked down studies of how often first-year college students made grammatical errors in freshman composition essays, going back nearly a century. She found that their error rate has barely risen at all. More astonishingly, today's freshman-comp essays are over six times longer than they were back then, and also generally more complex. "Student essayists of the early twentieth century often wrote essays on set topics like 'spring flowers,'" Lunsford tells me, "while those in the 1980s most often wrote personal experience narratives. Today's students are much more likely to write essays that present an argument, often with evidence to back them up"—a much more challenging task. And as for all those benighted texting short forms, like LOL, that have supposedly metastasized in young people's formal writing? Mostly nonexistent. "Our findings do not support such fears," Lunsford wrote in a paper describing her research, adding, "In fact, we found almost no instances of IM terms." Other studies have generally backed up Lunsford's observations: one analyzed 1.5 million words from instant messages by teens and found that even there, only 3 percent of the words used were IM-style short forms. (And while spelling and capitalization could be erratic, not all was awry; for example, youth substituted "u" for "you" only 8.6 percent of the time they wrote the word.) Others have found that kids who message a lot appear to have slightly better spelling and

literacy abilities than those who don't. At worst, messaging—with its half-textual, half-verbal qualities—might be reinforcing a preexisting social trend toward people writing more casually in otherwise formal situations, like school essays or the workplace.

In 2001, Lunsford got interested in the writing her students were doing everywhere—not just in the classroom, but outside it. She began the five-year Stanford Study of Writing, and she convinced 189 students to give her copies of everything they wrote, all year long, in any format: class papers, memos, e-mails, blog and discussion-board posts, text messages, instant-message chats, and more. Five years later, she'd collected nearly fifteen thousand pieces of writing and discovered something notable: The amount of writing kids did outside the class was huge. In fact, roughly 40 percent of everything they wrote was for pleasure, leisure, or socializing. "They're writing so much more than students before them ever did," she tells me. "It's stunning."

Lunsford also finds it striking how having an audience changed the students' writing outside the classroom. Because they were often writing for other people—the folks they were e-mailing with or talking with on a discussion board—they were adept at reading the tempo of a thread, adapting their writing to people's reactions. For Lunsford, the writing strategies of today's students have a lot in common with the Greek ideal of being a smart rhetorician: knowing how to debate, to marshal evidence, to listen to others, and to concede points. Their writing was constantly in dialogue with others.

"I think we are in the midst of a literacy revolution the likes of which we have not seen since Greek civilization," Lunsford tells me. The Greek oral period was defined by knowledge that was formed face-to-face, in debate with others. Today's online writing is like a merging of that culture and the Gutenberg print one. We're doing

more jousting that takes place in text but is closer in pacing to a face-to-face conversation. No sooner does someone assert something than the audience is reacting—agreeing, challenging, hysterically criticizing, flattering, or being abusive.

The upshot is that public thinking is often less about product than *process*. A newspaper runs a story, a friend posts a link on Facebook, a blogger writes a post, and it's interesting. But the real intellectual action often takes place in the comments. In the spring of 2011, a young student at Rutgers University in New Jersey was convicted of using his webcam to spy on a gay roommate, who later committed suicide. It was a controversial case and a controversial verdict, and when the *New York Times* wrote about it, it ran a comprehensive story more than 1,300 words long. But the readers' comments were many times larger—1,269 of them, many of which were remarkably nuanced, replete with complex legal and ethical arguments. I learned considerably more about the Rutgers case in a riveting half hour of reading *New York Times* readers debate the case than I learned from the article, because the article—substantial as it was—could represent only a small number of facets of a terrifically complex subject.

Socrates might be pleased. Back when he was alive, twenty-five hundred years ago, society had begun shifting gradually from an oral mode to a written one. For Socrates, the advent of writing was dangerous. He worried that text was too inert: once you wrote something down, that text couldn't adapt to its audience. People would read your book and think of a problem in your argument or want clarifications of your points, but they'd be out of luck. For Socrates, this was deadly to the quality of thought, because in the Greek intellectual tradition, knowledge was formed in the cut and thrust of debate. In Plato's *Phaedrus*, Socrates outlines these fears:

I cannot help feeling, Phaedrus, that writing is unfortunately like painting; for the creations of the painter have the attitude of life, and yet if you ask them a question they preserve a solemn silence. And the same may be said of speeches. You would imagine that they had intelligence, but if you want to know anything and put a question to one of them, the speaker always gives one unvarying answer. And when they have been once written down they are tumbled about anywhere among those who may or may not understand them, and know not to whom they should reply, to whom not: and, if they are maltreated or abused, they have no parent to protect them; and they cannot protect or defend themselves.

Today's online writing meets Socrates halfway. It's *printish*, but with a roiling culture of oral debate attached. Once something interesting or provocative is published—from a newspaper article to a book review to a tweet to a photo—the conversation begins, and goes on, often ad infinitum, and even the original authors can dive in to defend and extend their writing.

The truth is, of course, that knowledge has always been created via conversation, argument, and consensus. It's just that for the last century of industrial-age publishing, that process was mostly hidden from view. When I write a feature for a traditional print publication like *Wired* or *The New York Times*, it involves scores of conversations, conducted through e-mail and on the phone. The editors and I have to agree upon what the article will be about; as they edit the completed piece, the editors and fact-checkers will fix mistakes and we'll debate whether my paraphrase of an interviewee's point of view is too terse or glib. By the time we're done, we'll have generated a conversation about the article that's at least as long as the article itself (and probably far longer if you transcribed

our phone calls). The same thing happens with every book, documentary, or scientific paper—but because we don't see the sausage being made, we in the audience often forget that most information is forged in debate. I often wish traditional publishers let their audience see the process. I suspect readers would be intrigued by how magazine fact-checkers improve my columns by challenging me on points of fact, and they'd understand more about why material gets left out of a piece—or left in it.

Wikipedia has already largely moved past its period of deep suspicion, when most academics and journalists regarded it as utterly untrustworthy. Ever since the 2005 story in *Nature* that found Wikipedia and the *Encyclopedia Britannica* to have fairly similar error rates (four errors per article versus three, respectively), many critics now grudgingly accept Wikipedia as “a great place to start your research, and the worst place to end it.” Wikipedia's reliability varies heavily across the site, of course. Generally, articles with large and active communities of contributors are more accurate and complete than more marginal ones. And quality varies by subject matter; a study commissioned by the Wikipedia Foundation itself found that in the social sciences and humanities, the site is 10 to 16 percent less accurate than some expert sources.

But as the author David Weinberger points out, the deeper value of Wikipedia is that it makes transparent the arguments that go into the creation of any article: click on the “talk” page and you'll see the passionate, erudite conversations between Wikipedians as they hash out an item. Wikipedia's process, Weinberger points out, is a part of its product, arguably an indispensable part. Whereas the authority of traditional publishing relies on expertise—*trust us because our authors are vetted by our experience, their credentials, or the marketplace*—conversational media gains authority by revealing its mechanics. James Bridle, a British writer, artist, and publisher, made this point neatly when he took the entire text of every edit of

Wikipedia's much-disputed entry on the Iraq War during a five-year period and printed it as a set of twelve hardcover books. At nearly seven thousand pages, it was as long as an encyclopedia itself. The point, Bridle wrote, was to make visible just how much debate goes into the creation of a factual record: “This is historiography. This is what culture actually looks like: a process of argument, of dissenting and accreting opinion, of gradual and not always correct codification.” Public thinking is messy, but so is knowledge.

I'm not suggesting here, as have some digital utopians (and dystopians), that traditional “expert” forms of thinking and publishing are obsolete, and that expertise will corrode as the howling hive mind takes over. Quite the opposite. I work in print journalism, and now in print books, because the “typographical fixity” of paper—to use Elizabeth Eisenstein's lovely phrase—is a superb tool for focusing the mind. Constraints can impose creativity and rigor. When I have only six hundred words in a magazine column to make my point, I'm forced to make decisions about what I'm willing to commit to print. Slowing down also gives you time to consult a ton of sources and intuit hopefully interesting connections among them. The sheer glacial nature of the enterprise—spending years researching a book and writing it—is a cognitive strength, a gift that industrial processes gave to civilization. It helps one escape the speed loop of the digital conversation, where it's easy to fall prey to what psychologists call recency: Whatever's happening *right now* feels like the most memorable thing, so responding right now feels even *more* urgent. (This is a problem borrowed from face-to-face conversation: You won't find a lot of half-hour-long, thoughtful pauses in coffeehouse debates either.) And while traditional “expert” media are going to evolve in form and style, I doubt they're going to vanish, contrary to some of the current hand-wringing and gloating over that prospect. Business models for traditional reportage might be foundering, but interest is not: one analysis by HP Labs looked

at Twitter's "trending topics" and found that a majority of the most retweeted sources were mainstream news organizations like CNN, *The New York Times*, and Reuters.

The truth is that old and new modes of thinking aren't mutually exclusive. Knowing when to shift between public and private thinking—when to blast an idea online, when to let it slow bake—is a crucial new skill: cognitive diversity. When I get blocked while typing away at a project on my computer, I grab a pencil and paper, so I can use a tactile, swoopy, this-connects-to-that style of writing to unclog my brain. Once an idea is really flowing on paper, I often need to shift to the computer, so my seventy-words-per-minute typing and on-tap Google access can help me move swiftly before I lose my train of thought.

Artificial intelligence pioneer Marvin Minsky describes human smarts as stemming from the various ways our brains will tackle a problem; we'll simultaneously throw logic, emotion, metaphor, and crazy associative thinking at it. This works with artificial thinking tools, too. Spent too much time babbling online? Go find a quiet corner and read. Spent a ton of time working quietly alone? Go bang your ideas against other people online.

Ethan Hein is a musician who lives not far from me in Brooklyn. He teaches music and produces songs and soundtracks for indie movies and off-Broadway shows.

But most people know him as a guy who answers questions.

Tons of them. From strangers.

Hein is an enthusiastic poster on Quora, one of the current crop of question-answering sites: anyone can show up and ask a question, and anyone can answer. Hein had long been an online extrovert, blogging about music and tweeting. But he could also be, like many of us, lazy about writing. "I was always a half-assed journal

keeper," he tells me. "It was like, I should write something—wait a minute, what's on TV?" But in early 2011 he stumbled upon Quora and found the questions perversely stimulating. (Question: "What does the human brain find exciting about syncopated rhythm and breakbeats?" Hein's answer began: "Predictable unpredictability. The brain is a pattern-recognition machine . . .") Other times, he chimed in on everything from neuroscience and atheism to "What is it like to sleep in the middle of a forest?" (A: "Sleeping in the woods gratifies our biophilia.") Within a year, he was hooked.

"I will happily shuffle through the unanswered questions as a form of *entertainment*," Hein says. "My wife is kind of worried about me. But I'm like, 'Look, I'd be using this time to play *World of Warcraft*. And this is better—this is contributing. To the world!'" He even found that answering questions on Quora invigorated his blogging, because once he'd researched a question and pounded out a few paragraphs, he could use the answer as the seed for a new post. In barely one year he'd answered over twelve hundred questions and written about ninety thousand words. I tell him that's the length of a good-sized nonfiction hardcover book, and, as with Ory Okolloh and her two telephone books' worth of online writing, he seems stunned.

Public thinking is powerful, but it's hard to do. It's *work*. Sure, you get the good—catalyzing multiples, learning from the feedback. But it can be exhausting. Digital tools aren't magical pixie dust that makes you smarter. The opposite is true: they give up the rewards only if you work hard and master them, just like the cognitive tools of previous generations.

But as it turns out, there are structures that can make public thinking easier—and even irresistible.

Question answering is a powerful example. In the 1990s, question-answering sites like Answerbag.com began to emerge; by now there are scores of them. The sheer volume of questions answered is

remarkable: over one billion questions have been answered at the English version of Yahoo Answers, with one study finding the average answerer has written about fifty-one replies. In Korea, the search engine Naver set up shop in 1999 but realized there weren't very many Korean-language Web sites in existence, so it set up a question-answering forum, which became one of its core offerings. (And since all those questions are hosted in a proprietary database that Google can't access, Naver has effectively sealed Google out from the country, a neat trick.) Not all the answers, or questions, are good; Yahoo Answers in particular has become the butt of jokes for hosting spectacularly illiterate queries ("I CAN SMELL EVERYTHING MASSIVE HEAD ACHE?") or math students posting homework questions, hoping they'll be answered. (They usually are.) But some, like Quora, are known for cultivating thought-provoking questions and well-written answers. One of my favorite questions was "Who is history's greatest badass, and why?"—which provoked a twenty-two-thousand-word rush of answers, one of which described former U.S. president Theodore Roosevelt being shot by a would-be assassin before a speech and then, bleeding profusely, continuing to give the 1.5-hour-long address.

Why do question sites produce such outpourings of answers? It's because the format is a clever way of encouraging people to formalize and share knowledge. People walk around with tons of information and wisdom in their heads but with few outlets to show it off. Having your own Web site is powerful, but comparatively few people are willing to do the work. They face the blank-page problem. What should I say? Who *cares* what I say? In contrast, when you see someone asking a question on a subject you know about, it catalyzes your desire to speak up.

"Questions are a really useful service for curing writer's block," as Charlie Cheever, the soft-spoken cofounder of Quora, tells me. "You might think you want to start a blog, but you wind up being

afraid to write a blog post because there's this sense of, who asked *you?*" Question answering provides a built-in, instant audience of at least one—the original asker. This is another legacy of Plato's Socratic dialogues, in which Socrates asks questions of his debating partners (often faux-naive, concern-trolling ones, of course) and they pose questions of him in turn. Web authors long ago turned this into a literary form that has blossomed: the FAQ, a set of mock-Socratic questions authors pose to themselves as a way of organizing information.

It's an addictive habit, apparently. Academic research into question-answering sites has found that answering begets answering: people who respond to questions are likely to stick around for months and answer even more. Many question-answering sites have a psychological architecture of rewards, such as the ability of members to give positive votes (or award "points") for good answers. But these incentives may be secondary to people's altruism and the sheer joy of helping people out, as one interview survey of Naver users discovered. The Naver users said that once they stumbled across a question that catalyzed their expertise, they were hooked; they couldn't help responding. "Since I was a doctor, I was browsing the medical directories. I found a lot of wrong answers and information and was afraid they would cause problems," as one Naver contributor said. "So I thought I'd contribute in fixing it, hoping that it'd be good for the society." Others found that the act of writing answers helped organize their own thoughts—the generation effect in a nutshell. "My first intention [in answering] was to organize and review my knowledge and practice it by explaining it to others," one explained.

These sites have formalized question answering as a vehicle for public thinking, but they didn't invent it. In almost any online community, answering questions frequently forms the backbone of conversation, evolving on a grassroots level. Several years ago while

reading YouBeMom, an anonymous forum for mothers, I noticed that users had created a clever inversion of the question-answering format: a user would post a description of their job and ask if anyone had questions. The ploy worked in both directions, encouraging people to ask questions they might never have had the opportunity to ask. The post “ER nurse here—questions?” turned into a sprawling discussion, hundreds of postings long, about the nurse’s bloodiest accidents, why gunshot attacks were decreasing, and how ballooning ER costs are destroying hospital budgets. (An even more spellbinding conversation emerged the night a former prostitute opened up the floor for questions.) Though it’s hard to say where it emerged, the “I am a . . .” format has become, like the FAQ, another literary genre the Internet has ushered into being; on the massive discussion board Reddit, there are dozens of “IAMA” threads started each day by everyone from the famous (the comedian Louis C.K., Barack Obama) to people with intriguing experiences (“IAMA Female Vietnam Veteran”; “IAMA former meth lab operator”; “IAMA close friend of Charlie Sheen since 1985”).

I’m focusing on question answering, but what’s really at work here is what publisher and technology thinker Tim O’Reilly calls the “architecture of participation.” The future of public thinking hinges on our ability to create tools that bring out our best: that encourage us to organize our thoughts, create audiences, make connections. Different forms encourage different styles of talk.

Microblogging created a torrent of public thinking by making a virtue of its limits. By allowing people to write only 140 characters at a time, Twitter neatly routed around the “blank page” problem: everybody can think of at least *that* many words to say. Facebook provoked a flood of writing by giving users audiences composed of people they already knew well from the offline world, people they knew cared about what they had to say. Texting offered a style of conversation that was more convenient than voice calls (and cheaper,

in developing countries), and the asynchronicity created pauses useful for gathering your thoughts (or waiting until your boss’s back was turned so you could sneak in a conversation). One size doesn’t fit all, cognitively speaking. I know people who engage in arguments about music or politics with friends on Facebook because it’s an extension of offline contact, while others find the presence of friends claustrophobic; they find it more freeing and stimulating to talk with comparative strangers on open-ended discussion boards.



Clearly, public speech can be enormously valuable. But what about stuff that isn’t? What about the *repellent* public speech? When you give everyday people the ability to communicate, you release not just brilliant bons mots and incisive conversations, but also ad hominem attacks, fury, and “trolls”—people who jump into discussion threads solely to destabilize them. The combination of distance and pseudonymity (or sometimes total anonymity) can unlock people’s worst behavior, giving them license to say brutal things they’d never say to someone’s face.

This abuse isn’t evenly distributed. It’s much less often directed at men, particularly white men like me. In contrast, many women I know—probably most—find that being public online inevitably attracts a wave of comments, ranging from dismissal to assessments of their appearance to flat-out rape threats. This is particularly true if they’re talking about anything controversial or political. Or even intellectual: “An opinion, it seems, is the short skirt of the Internet,” as Laurie Penny, a British political writer, puts it. This abuse is also heaped on blacks and other minorities in the United States, or any subordinated group. Even across lines of party politics, discussion threads quickly turn toxic in highly personal ways.

How do we end this type of abuse? Alas, we probably can’t, at least not completely—after all, this venom is rooted in real-world

biases that go back centuries. The Internet didn't create these prejudices; it gave them a new stage.

But there are, it turns out, techniques to curtail online abuse, sometimes dramatically. In fact, some innovators are divining, through long experience and experimentation, key ways of managing conversation online—not only keeping it from going septic, but improving it.

Consider the example of Ta-Nahesi Coates. Coates is a senior editor at *The Atlantic Monthly*, a magazine of politics and culture; he ran a personal blog for years and moved it over to the *Atlantic* five years ago. Coates posts daily on a dizzying array of subjects: movies, politics, economic disparities, the Civil War, TV shows, favorite snippets of poetry, or whether pro football is too dangerous to play. Coates, who is African American, is also well known as an eloquent and incisive writer on race, and he posts about that frequently. Yet his forum is amazingly abuse-free: comments spill into the hundreds without going off the rails. “This is the most hot-button issue in America, and folks have managed to keep a fairly level head,” he tells me.

The secret is the work Coates puts into his discussion board. Before he was a blogger himself, he'd noticed the terrible comments at his favorite political blogs, like that of Matt Yglesias. “Matt could be talking about parking and urban issues, and he'd have ten comments, and somebody would invariably say something racist.” Coates realized that negative comments create a loop: they poison the atmosphere, chasing off productive posters.

So when he started his own personal blog, he decided to break that loop. The instant he saw something abusive, he'd delete it, banning repeat offenders. Meanwhile, he went out of his way to encourage the smart folks, responding to them personally and publicly, so they'd be encouraged to stay and talk. And Coates was unfaill-

ingly polite and civil himself, to help set community standards. Soon several dozen regular commenters emerged, and they got to know each other, talking as much to each other as to Coates. (They've even formed their own Facebook group and have held “meet-ups.”) Their cohesion helped cement the culture of civility even more; any troll today who looks at the threads can quickly tell this community isn't going to tolerate nastiness. The *Atlantic* also deploys software that lets users give an “up” vote to the best comments, which further helps reinforce quality. Given that the community has good standards, the first comment thread you'll see at the bottom of a Coates post is likely to be the cleverest—and not, as at sites that don't manage their comments and run things chronologically, the first or last troll to have stopped by.

This is not to say it's a love fest or devoid of conflict. The crowd argues heatedly and often takes Coates to task for his thinking; he cites their feedback in his own posts. “Being a writer does not mean you are smarter than everyone else. I *learn* things from these people,” he notes. But the debate transpires civilly and without name-calling. These days, Coates still tends the comments and monitors them but rarely needs to ban anyone. “It's much easier,” he adds.

What exactly do you call what Coates is doing, this mix of persuasion, listening, and good hosting, like someone skillfully tending bar? A few years ago, three Internet writers and thinkers—Deb Schultz, Heather Gold, and Kevin Mark—brainstormed on what to label it. On the suggestion of Theresa Nielsen Hayden, a longtime host of online communities, they settled on a clever term: “tummeling,” derived from the Yiddish *tummler*, the person at a party responsible for keeping the crowd engaged and getting them dancing at a wedding. Tummlers are the social adepts of online conversation. “They're catalysts and bridge builders,” Schultz tells me. “It's not about technology. It's about the human factor.” They know how to

be empathetic, how to draw people out: “A good tummler reads the room,” Gold adds. “Quieter people have a disproportionately strong impact on conversational flow when drawn out and heard.”

Look behind any high-functioning discussion forum online and you’ll find someone doing tummeling. Without it, you get chaos. That’s why YouTube is a comment cesspool; there is no culture of moderating comments. It’s why you frequently see newspaper Web pages filled with toxic comments. They haven’t assigned anyone to be the tummler.

Newspapers and YouTube also have another problem, which is that they’re always trying to get bigger. But as Coates and others have found, conversation works best when it’s smaller. Only in a more tightly knit group can participants know each other. Newspapers, in contrast, work under the advertising logic of “more is better.” This produces unfocused, ad hoc, drive-by audiences that can never be corralled into community standards. Coates jokes about going to a major U.S. newspaper and seeing a link to the discussion threads—*Come on in! We have 2,000 comments! “That’s a bar I don’t want to go into! They don’t have any security!”* he says. These sites are trying for scale—but conversation doesn’t scale.

There are other tools emerging to help manage threads, such as requiring real name identity, as with Facebook comments; removing anonymity can bring in accountability, since people are less likely to be abusive if their actual name is attached to the abuse. Mind you, Coates isn’t opposed *per se* to anonymity or to crazy, free-range places like Reddit. “Those environments catalyze a lot of rancor, sure, but also candor. The fact that places like that exist might make it even easier to do what I do,” he notes.

Tummeling isn’t a total solution. It works only when you control the space and can kick out undesirables. Services like Twitter are more open and thus less manageable. But even in those spaces, tummeling is a digital-age skill that we will increasingly need to learn,

even formally *teach*; if this aspect of modern civics became widespread enough, it could help reform more and more public spaces online. There’s a pessimistic view, too. You could argue that the first two decades of open speech have set dreadful global standards and that the downsides of requiring targeted groups—say, young women—to navigate so much hate online aren’t worth the upsides of public speech. That’s a reasonable caveat. When it comes to public thinking, you need to accept the bad with the good, but there’s a lot of bad to accept.

What tools will create new forms of public thinking in the years to come? With mobile phones, our personal geography is becoming newly relevant in a new way. GPS turns your location into a fresh source of multiples, because it can figure out if there are other people nearby sharing your experience (say, at a concert or a park). An early success of this kind was Grindr, a phone app that lets gay men broadcast their location and status messages and locate other gay men nearby (proving again the technology truism that sex and pornography are always at the forefront of tech innovation).

The ability of phones to broadcast their location has even weirder effects, because it can turn geography into a message board, with apps that embed conversations in specific physical spaces. For example, when the Occupy Wall Street movement flared in New York City, some of the activists began using a mobile app called Vibe that let them post anonymous messages that were tagged to physical locations around Wall Street: they’d discuss where police were about to crack down or leave notes describing events they’d seen. This is bleeding into everyday life, with services that let people embed photos and thoughts on maps and engage in location-based conversations. It’s the first stage of conversational “augmented reality”: public thinking woven into our real-world public space.

I also suspect that as more forms of media become digital, they'll become sites for public thinking—particularly digital books. Books have always propelled smart conversations; the historic, face-to-face book club has migrated rapidly online, joining the sprawling comments at sites like Goodreads. But the pages of e-books are themselves likely to become the sites of conversations. Already readers of many e-books—on the Kindle, the Nook, and other e-readers—share comments and highlights. Marginalia may become a new type of public thinking, with the smartest remarks from other readers becoming part of how we make sense of a book. (Bob Stein, head of the Institute for the Future of the Book, imagines a cadre of marginaliaists becoming so well liked that people pay to read their mark-ups.) The truth is, whatever new digital tools come around, curious people are going to colonize them. We're social creatures, so we think socially.

But there's one interesting kink. For most of this chapter I've been talking about one type of publishing—writing in text. It's one of our oldest and most robust tools for recording and manipulating ideas. But the digital age is also producing a Cambrian explosion in different media that we're using to talk, and think, with each other—including images, video, and data visualization. The difference is, while we're taught in school how to write and read, our traditional literacy focuses less on these new modes of publishing. We're working them out on our own for now and discovering just how powerful they can be.



The New Literacies

How do you tackle a problem that affects the fabric of democracy but also happens to be, well, boring?

Ask Costas Panagopoulos. A professor of political science at Fordham University in New York, Panagopoulos is an expert on gerrymandering, the tawdry two-hundred-year-old political phenomenon by which politicians redraw the boundaries of their districts in order to exclude anyone who won't vote for them. In theory, redistricting isn't harmful; indeed, laws require the regular rejiggering of maps to make sure that as the population shifts, it's adequately represented. But in practice, politicians manipulate this process in order to cement their own power. In the United States, Democrats try to herd liberal urbanites and blacks into their districts' boundaries while pushing out gun-loving rural folk. Republicans do the reverse.

Politicians worldwide love this trick, but in New York State they've made it an art form. In the last fifty years, they've redrawn their electoral districts in such nakedly self-serving ways it's a statewide joke. (One district was redrawn so tortuously, a voting-rights advocate said that it looked like "Abraham Lincoln riding on a vacuum cleaner.") The result is a rigid, unchanging political terrain: once someone's in office, they almost can't be voted out. From 2002 to 2010, a slender 4 percent of incumbent New York state politicians lost. In 2010, nearly one out of five politicians didn't even