

The I in the Internet

In the beginning the internet seemed good. "I was in love with the internet the first time I used it at my dad's office and thought it was the ULTIMATE COOL," I wrote, when I was ten, on an Angelfire subpage titled "The Story of How Jia Got Her Web Addiction." In a text box superimposed on a hideous violet background, I continued:

But that was in third grade and all I was doing was going to Beanie Baby sites. Having an old, icky bicky computer at home, we didn't have the Internet. Even AOL seemed like a far-off dream. Then we got a new top-o'-the-line computer in spring break '99, and of course it came with all that demo stuff. So I finally had AOL and I was completely amazed at the marvel of having a profile and chatting and IMS!!

Then, I wrote, I discovered personal webpages. ("I was astonished!") I learned HTML and "little Javascript trickies." I built my own site on the beginner-hosting site Expage, choosing pastel colors and then switching to a "starry night theme." Then I ran out of space, so I "decided to move to Angelfire. Wow." I learned how to make my own graphics. "This was all in the course of four months,"

I wrote, marveling at how quickly my ten-year-old internet citizenry was evolving. I had recently revisited the sites that had once inspired me, and realized “how much of an idiot I was to be wowed by *that*.”

I have no memory of inadvertently starting this essay two decades ago, or of making this Angelfire subpage, which I found while hunting for early traces of myself on the internet. It’s now eroded to its skeleton: its landing page, titled “THE VERY BEST,” features a sepia-toned photo of Andie from *Dawson’s Creek* and a dead link to a new site called “THE FROSTED FIELD,” which is “BETTER!” There’s a page dedicated to a blinking mouse GIF named Susie, and a “Cool Lyrics Page” with a scrolling banner and the lyrics to Smash Mouth’s “All Star,” Shania Twain’s “Man! I Feel Like a Woman!” and the TLC diss track “No Pigeons,” by Sporty Thievez. On an FAQ page—there was an FAQ page—I write that I had to close down my customizable cartoon-doll section, as “the response has been enormous.”

It appears that I built and used this Angelfire site over just a few months in 1999, immediately after my parents got a computer. My insane FAQ page specifies that the site was started in June, and a page titled “Journal”—which proclaims, “I am going to be completely honest about my life, although I won’t go too deeply into personal thoughts, though”—features entries only from October. One entry begins: “It’s so HOT outside and I can’t count the times acorns have fallen on my head, maybe from exhaustion.” Later on, I write, rather prophetically: “I’m going insane! I literally am addicted to the web!”

In 1999, it felt different to spend all day on the internet. This was true for everyone, not just for ten-year-olds: this was the *You’ve Got Mail* era, when it seemed that the very worst thing that could happen online was that you might fall in love with your business rival. Throughout the eighties and nineties, people had been gathering on the internet in open forums, drawn, like butterflies, to the puddles and blossoms of other people’s curiosity

and expertise. Self-regulated newsgroups like Usenet cultivated lively and relatively civil discussion about space exploration, meteorology, recipes, rare albums. Users gave advice, answered questions, made friendships, and wondered what this new internet would become.

Because there were so few search engines and no centralized social platforms, discovery on the early internet took place mainly in private, and pleasure existed as its own solitary reward. A 1995 book called *You Can Surf the Net!* listed sites where you could read movie reviews or learn about martial arts. It urged readers to follow basic etiquette (don’t use all caps; don’t waste other people’s expensive bandwidth with overly long posts) and encouraged them to feel comfortable in this new world (“Don’t worry,” the author advised. “You have to *really* mess up to get flamed.”). Around this time, GeoCities began offering personal website hosting for dads who wanted to put up their own golfing sites or kids who built glittery, blinking shrines to Tolkien or Ricky Martin or unicorns, most capped off with a primitive guest book and a green-and-black visitor counter. GeoCities, like the internet itself, was clumsy, ugly, only half functional, and organized into neighborhoods: /area51/ was for sci-fi, /westhollywood/ for LGBTQ life, /enchantedforest/ for children, /petsburgh/ for pets. If you left GeoCities, you could walk around other streets in this ever-expanding village of curiosities. You could stroll through Expage or Angelfire, as I did, and pause on the thoroughfare where the tiny cartoon hamsters danced. There was an emergent aesthetic—blinking text, crude animation. If you found something you liked, if you wanted to spend more time in any of these neighborhoods, you could build your own house from HTML frames and start decorating.

This period of the internet has been labeled Web 1.0—a name that works backward from the term Web 2.0, which was coined by the writer and user-experience designer Darcy DiNucci in an article called “Fragmented Future,” published in 1999. “The Web

we know now," she wrote, "which loads into a browser window in essentially static screenfuls, is only an embryo of the Web to come. The first glimmerings of Web 2.0 are beginning to appear. . . . The Web will be understood not as screenfuls of texts and graphics but as a transport mechanism, the ether through which interactivity happens." On Web 2.0, the structures would be dynamic, she predicted: instead of houses, websites would be portals, through which an ever-changing stream of activity—status updates, photos—could be displayed. What you did on the internet would become intertwined with what everyone else did, and the things other people liked would become the things that you would see. Web 2.0 platforms like Blogger and Myspace made it possible for people who had merely been taking in the sights to start generating their own personalized and constantly changing scenery. As more people began to register their existence digitally, a pastime turned into an imperative: you had to register yourself digitally to exist.

In a *New Yorker* piece from November 2000, Rebecca Mead profiled Meg Hourihan, an early blogger who went by Megnut. In just the prior eighteen months, Mead observed, the number of "weblogs" had gone from fifty to several thousand, and blogs like Megnut were drawing thousands of visitors per day. This new internet was social ("a blog consists primarily of links to other Web sites and commentary about those links") in a way that centered on individual identity (Megnut's readers knew that she wished there were better fish tacos in San Francisco, and that she was a feminist, and that she was close with her mom). The blogosphere was also full of mutual transactions, which tended to echo and escalate. The "main audience for blogs is other bloggers," Mead wrote. Etiquette required that, "if someone blogs your blog, you blog his blog back."

Through the emergence of blogging, personal lives were becoming public domain, and social incentives—to be liked, to be seen—were becoming economic ones. The mechanisms of inter-

net exposure began to seem like a viable foundation for a career. Hourihan cofounded Blogger with Evan Williams, who later cofounded Twitter. JenniCam, founded in 1996 when the college student Jennifer Ringley started broadcasting webcam photos from her dorm room, attracted at one point up to four million daily visitors, some of whom paid a subscription fee for quicker-loading images. The internet, in promising a potentially unlimited audience, began to seem like the natural home of self-expression. In one blog post, Megnut's boyfriend, the blogger Jason Kottke, asked himself why he didn't just write his thoughts down in private. "Somehow, that seems strange to me though," he wrote. "The Web is the place for you to express your thoughts and feelings and such. To put those things elsewhere seems absurd."

Every day, more people agreed with him. The call of self-expression turned the village of the internet into a city, which expanded at time-lapse speed, social connections bristling like neurons in every direction. At ten, I was clicking around a web ring to check out other Angelfire sites full of animal GIFs and Smash Mouth trivia. At twelve, I was writing five hundred words a day on a public LiveJournal. At fifteen, I was uploading photos of myself in a miniskirt on Myspace. By twenty-five, my job was to write things that would attract, ideally, a hundred thousand strangers per post. Now I'm thirty, and most of my life is inextricable from the internet, and its mazes of incessant forced connection—this feverish, electric, unlivable hell.

As with the transition between Web 1.0 and Web 2.0, the curdling of the social internet happened slowly and then all at once. The tipping point, I'd guess, was around 2012. People were losing excitement about the internet, starting to articulate a set of new truisms. Facebook had become tedious, trivial, exhausting. Instagram seemed better, but would soon reveal its underlying function as a three-ring circus of happiness and popularity and success. Twitter, for all its discursive promise, was where everyone tweeted complaints at airlines and bitched about articles that had been

commissioned to make people bitch. The dream of a better, truer self on the internet was slipping away. Where we had once been free to be ourselves online, we were now *chained* to ourselves online, and this made us self-conscious. Platforms that promised connection began inducing mass alienation. The freedom promised by the internet started to seem like something whose greatest potential lay in the realm of misuse.

Even as we became increasingly sad and ugly on the internet, the mirage of the better online self continued to glimmer. As a medium, the internet is defined by a built-in performance incentive. In real life, you can walk around living life and be visible to other people. But you can't just walk around and be visible on the internet—for anyone to see you, you have to *act*. You have to communicate in order to maintain an internet presence. And, because the internet's central platforms are built around personal profiles, it can seem—first at a mechanical level, and later on as an encoded instinct—like the main purpose of this communication is to make yourself look good. Online reward mechanisms beg to substitute for offline ones, and then overtake them. This is why everyone tries to look so hot and well-traveled on Instagram; this is why everyone seems so smug and triumphant on Facebook; this is why, on Twitter, making a righteous political statement has come to seem, for many people, like a political good in itself.

This practice is often called “virtue signaling,” a term most often used by conservatives criticizing the left. But virtue signaling is a bipartisan, even apolitical action. Twitter is overrun with dramatic pledges of allegiance to the Second Amendment that function as intra-right virtue signaling, and it can be something like virtue signaling when people post the suicide hotline after a celebrity death. Few of us are totally immune to the practice, as it intersects with a real desire for political integrity. Posting photos from a protest against border family separation, as I did while writing this, is a microscopically meaningful action, an expression

of genuine principle, and also, inescapably, some sort of attempt to signal that I am good.

Taken to its extreme, virtue signaling has driven people on the left to some truly unhinged behavior. A legendary case occurred in June 2016, after a two-year-old was killed at a Disney resort—dragged off by an alligator while playing in a no-swimming-allowed lagoon. A woman, who had accumulated ten thousand Twitter followers with her posts about social justice, saw an opportunity and tweeted, magnificently, “I’m so finished with white men’s entitlement lately that I’m really not sad about a 2yo being eaten by a gator because his daddy ignored signs.” (She was then pilloried by people who chose to demonstrate their own moral superiority through mockery—as I am doing here, too.) A similar tweet made the rounds in early 2018 after a sweet story went viral: a large white seabird named Nigel had died next to the concrete decoy bird to whom he had devoted himself for years. An outraged writer tweeted, “Even concrete birds do not owe you affection, Nigel,” and wrote a long Facebook post arguing that Nigel’s courtship of the fake bird exemplified . . . *rape culture*. “I’m available to write the feminist perspective on Nigel the gannet’s non-tragic death should anyone wish to pay me,” she added, underneath the original tweet, which received more than a thousand likes. These deranged takes, and their unnerving proximity to online monetization, are case studies in the way that our world—digitally mediated, utterly consumed by capitalism—makes communication about morality very easy but makes actual moral living very hard. You don’t end up using a news story about a dead toddler as a peg for white entitlement without a society in which the discourse of righteousness occupies far more public attention than the conditions that necessitate righteousness in the first place.

On the right, the online performance of political identity has been even wilder. In 2017, the social-media-savvy youth conserva-

tive group Turning Point USA staged a protest at Kent State University featuring a student who put on a diaper to demonstrate that “safe spaces were for babies.” (It went viral, as intended, but not in the way TPUSA wanted—the protest was uniformly roasted, with one Twitter user slapping the logo of the porn site Brazzers on a photo of the diaper boy, and the Kent State TPUSA campus coordinator resigned.) It has also been infinitely more consequential, beginning in 2014, with a campaign that became a template for right-wing internet-political action, when a large group of young misogynists came together in the event now known as Gamergate.

The issue at hand was, ostensibly, a female game designer perceived to be sleeping with a journalist for favorable coverage. She, along with a set of feminist game critics and writers, received an onslaught of rape threats, death threats, and other forms of harassment, all concealed under the banner of free speech and “ethics in games journalism.” The Gamergaters—estimated by *Deadspin* to number around ten thousand people—would mostly deny this harassment, either parroting in bad faith or fooling themselves into believing the argument that Gamergate was actually about noble ideals. Gawker Media, *Deadspin*’s parent company, itself became a target, in part because of its own aggressive disdain toward the Gamergaters: the company lost seven figures in revenue after its advertisers were brought into the maelstrom.

In 2016, a similar fiasco made national news in Pizzagate, after a few rabid internet denizens decided they’d found coded messages about child sex slavery in the advertising of a pizza shop associated with Hillary Clinton’s campaign. This theory was disseminated all over the far-right internet, leading to an extended attack on DC’s Comet Ping Pong pizzeria and everyone associated with the restaurant—all in the name of combating pedophilia—that culminated in a man walking into Comet Ping Pong and firing a gun. (Later on, the same faction would jump to the defense of Roy Moore, the Republican nominee for the Senate who was

accused of sexually assaulting teenagers.) The over-woke left could only dream of this ability to weaponize a sense of righteousness. Even the militant antifascist movement, known as antifa, is routinely disowned by liberal centrists, despite the fact that the antifa movement is rooted in a long European tradition of Nazi resistance rather than a nascent constellation of radically paranoid message boards and YouTube channels. The worldview of the Gamergaters and Pizzagaters was actualized and to a large extent vindicated in the 2016 election—an event that strongly suggested that the worst things about the internet were now *determining*, rather than reflecting, the worst things about offline life.

Mass media always determines the shape of politics and culture. The Bush era is inextricable from the failures of cable news; the executive overreaches of the Obama years were obscured by the internet’s magnification of personality and performance; Trump’s rise to power is inseparable from the existence of social networks that must continually aggravate their users in order to continue making money. But lately I’ve been wondering how everything got so *intimately* terrible, and why, exactly, we keep playing along. How did a huge number of people begin spending the bulk of our disappearing free time in an openly torturous environment? How did the internet get so bad, so confining, so inescapably personal, so politically determinative—and why are all those questions asking the same thing?

I’ll admit that I’m not sure that this inquiry is even productive. The internet reminds us on a daily basis that it is not at all rewarding to become aware of problems that you have no reasonable hope of solving. And, more important, the internet already is what it is. It has already become the central organ of contemporary life. It has already rewired the brains of its users, returning us to a state of primitive hyperawareness and distraction while overloading us with much more sensory input than was ever possible in primitive times. It has already built an ecosystem that runs on exploiting attention and monetizing the self. Even if you avoid the

internet completely—my partner does: he thought #tbt meant “truth be told” for ages—you still live in the world that this internet has created, a world in which selfhood has become capitalism’s last natural resource, a world whose terms are set by centralized platforms that have deliberately established themselves as near-impossible to regulate or control.

The internet is also in large part inextricable from life’s pleasures: our friends, our families, our communities, our pursuits of happiness, and—sometimes, if we’re lucky—our work. In part out of a desire to preserve what’s worthwhile from the decay that surrounds it, I’ve been thinking about five intersecting problems: first, how the internet is built to distend our sense of identity; second, how it encourages us to overvalue our opinions; third, how it maximizes our sense of opposition; fourth, how it cheapens our understanding of solidarity; and, finally, how it destroys our sense of scale.

In 1959, the sociologist Erving Goffman laid out a theory of identity that revolved around playacting. In every human interaction, he wrote in *The Presentation of Self in Everyday Life*, a person must put on a sort of performance, create an impression for an audience. The performance might be calculated, as with the man at a job interview who’s practiced every answer; it might be unconscious, as with the man who’s gone on so many interviews that he naturally performs as expected; it might be automatic, as with the man who creates the correct impression primarily because he is an upper-middle-class white man with an MBA. A performer might be fully taken in by his own performance—he might actually believe that his biggest flaw is “perfectionism”—or he might know that his act is a sham. But no matter what, he’s performing. Even if he stops *trying* to perform, he still has an audience, his actions still create an effect. “All the world is not, of course, a stage,

but the crucial ways in which it isn’t are not easy to specify,” Goffman wrote.

To communicate an identity requires some degree of self-delusion. A performer, in order to be convincing, must conceal “the discreditable facts that he has had to learn about the performance; in everyday terms, there will be things he knows, or has known, that he will not be able to tell himself.” The interviewee, for example, avoids thinking about the fact that his biggest flaw actually involves drinking at the office. A friend sitting across from you at dinner, called to play therapist for your trivial romantic hang-ups, has to pretend to herself that she wouldn’t rather just go home and get in bed to read Barbara Pym. No audience has to be physically present for a performer to engage in this sort of selective concealment: a woman, home alone for the weekend, might scrub the baseboards and watch nature documentaries even though she’d rather trash the place, buy an eight ball, and have a Craigslist orgy. People often make faces, in private, in front of bathroom mirrors, to convince themselves of their own attractiveness. The “lively belief that an unseen audience is present,” Goffman writes, can have a significant effect.

Offline, there are forms of relief built into this process. Audiences change over—the performance you stage at a job interview is different from the one you stage at a restaurant later for a friend’s birthday, which is different from the one you stage for a partner at home. At home, you might feel as if you could stop performing altogether; within Goffman’s dramaturgical framework, you might feel as if you had made it backstage. Goffman observed that we need both an audience to witness our performances as well as a backstage area where we can relax, often in the company of “teammates” who had been performing alongside us. Think of coworkers at the bar after they’ve delivered a big sales pitch, or a bride and groom in their hotel room after the wedding reception: everyone may still be performing, but they

feel at ease, unguarded, alone. Ideally, the outside audience has believed the prior performance. The wedding guests think they've actually just seen a pair of flawless, blissful newlyweds, and the potential backers think they've met a group of geniuses who are going to make everyone very rich. "But this imputation—this self—is a product of a scene that comes off, and is not a cause of it," Goffman writes. The self is not a fixed, organic thing, but a dramatic effect that emerges from a performance. This effect can be believed or disbelieved at will.

Online—assuming you buy this framework—the system metastasizes into a wreck. The presentation of self in everyday internet still corresponds to Goffman's playacting metaphor: there are stages, there is an audience. But the internet adds a host of other, nightmarish metaphorical structures: the mirror, the echo, the panopticon. As we move about the internet, our personal data is tracked, recorded, and resold by a series of corporations—a regime of involuntary technological surveillance, which subconsciously decreases our resistance to the practice of *voluntary* self-surveillance on social media. If we think about buying something, it follows us around everywhere. We can, and probably do, limit our online activity to websites that further reinforce our own sense of identity, each of us reading things written for people just like us. On social media platforms, everything we see corresponds to our conscious choices and algorithmically guided preferences, and all news and culture and interpersonal interaction are filtered through the home base of the profile. The everyday madness perpetuated by the internet is the madness of this architecture, which positions personal identity as the center of the universe. It's as if we've been placed on a lookout that oversees the entire world and given a pair of binoculars that makes everything look like our own reflection. Through social media, many people have quickly come to view all new information as a sort of direct commentary on *who they are*.

... because it is profitable. As Tim Wu writes

in *The Attention Merchants*, commerce has been slowly permeating human existence—entering our city streets in the nineteenth century through billboards and posters, then our homes in the twentieth century through radio and TV. Now, in the twenty-first century, in what appears to be something of a final stage, commerce has filtered into our identities and relationships. We have generated billions of dollars for social media platforms through our desire—and then through a subsequent, escalating economic and cultural requirement—to replicate for the internet who we know, who we think we are, who we want to be.

Selfhood buckles under the weight of this commercial importance. In physical spaces, there's a limited audience and time span for every performance. Online, your audience can hypothetically keep expanding forever, and the performance never has to end. (You can essentially be on a job interview in perpetuity.) In real life, the success or failure of each individual performance often plays out in the form of concrete, physical action—you get invited over for dinner, or you lose the friendship, or you get the job. Online, performance is mostly arrested in the nebulous realm of sentiment, through an unbroken stream of hearts and likes and eyeballs, aggregated in numbers attached to your name. Worst of all, there's essentially no backstage on the internet; where the offline audience necessarily empties out and changes over, the online audience never has to leave. The version of you that posts memes and selfies for your pre-cal classmates might end up sparring with the Trump administration after a school shooting, as happened to the Parkland kids—some of whom became so famous that they will never be allowed to drop the veneer of performance again. The self that traded jokes with white supremacists on Twitter is the self that might get hired, and then fired, by *The New York Times*, as happened to Quinn Norton in 2018. (Or, in the case of Sarah Jeong, the self that made jokes *about* white people might get Gamergated after being hired at the *Times* a few months thereafter.) People who maintain a public internet profile

are building a self that can be viewed simultaneously by their mom, their boss, their potential future bosses, their eleven-year-old nephew, their past and future sex partners, their relatives who loathe their politics, as well as anyone who cares to look for any possible reason. Identity, according to Goffman, is a series of claims and promises. On the internet, a highly functional person is one who can promise everything to an indefinitely increasing audience at all times.

Incidents like Gamergate are partly a response to these conditions of hyper-visibility. The rise of trolling, and its ethos of disrespect and anonymity, has been so forceful in part because the internet's insistence on consistent, approval-worthy identity is so strong. In particular, the misogyny embedded in trolling reflects the way women—who, as John Berger wrote, have always been required to maintain an external awareness of their own identity—often navigate these online conditions so profitably. It's the self-calibration that I learned as a girl, as a woman, that has helped me capitalize on "having" to be online. My only experience of the world has been one in which personal appeal is paramount and self-exposure is encouraged; this legitimately unfortunate paradigm, inhabited first by women and now generalized to the entire internet, is what trolls loathe and actively repudiate. They destabilize an internet built on transparency and likability. They pull us back toward the chaotic and the unknown.

Of course, there are many better ways of making the argument against hyper-visibility than trolling. As Werner Herzog told *GQ*, in 2011, speaking about psychoanalysis: "We have to have our dark corners and the unexplained. We will become uninhabitable in a way an apartment will become uninhabitable if you illuminate every single dark corner and under the table and wherever—you cannot live in a house like this anymore."

The first time I was ever paid to publish anything, it was 2013, the end of the blog era. Trying to make a living as a writer with the internet as a standing precondition of my livelihood has given me some professional motivation to stay active on social media, making my work and personality and face and political leanings and dog photos into a continually updated record that anyone can see. In doing this, I have sometimes felt the same sort of unease that washed over me when I was a cheerleader and learned how to convincingly fake happiness at football games—the feeling of acting as if conditions are fun and normal and worthwhile in the hopes that they will just magically become so. To try to write online, more specifically, is to operate on a set of assumptions that are already dubious when limited to writers and even more questionable when turned into a categorical imperative for everyone on the internet: the assumption that speech has an impact, that it's something like action; the assumption that it's fine or helpful or even *ideal* to be constantly writing down what you think.

I have benefited, I mean, from the internet's unhealthy focus on opinion. This focus is rooted in the way the internet generally minimizes the need for physical action: you don't have to do much of anything but sit behind a screen to live an acceptable, possibly valorized, twenty-first-century life. The internet can feel like an astonishingly direct line to reality—click if you want something and it'll show up at your door two hours later; a series of tweets goes viral after a tragedy and soon there's a nationwide high school walkout—but it can also feel like a shunt diverting our energy *away* from action, leaving the real-world sphere to the people who already control it, keeping us busy figuring out the precisely correct way of explaining our lives. In the run-up to the 2016 election and increasingly so afterward, I started to feel that there was almost nothing I could do about ninety-five percent of the things I cared about other than form an opinion—and that the conditions that allowed me to live in mild everyday hysterics about an unlimited supply of terrible information were related to the con-

ditions that were, at the same time, consolidating power, sucking wealth upward, far outside my grasp.

I don't mean to be naïvely fatalistic, to act like *nothing* can be done about *anything*. People are making the world better through concrete footwork every day. (Not me—I'm too busy sitting in front of the internet!) But their time and labor, too, has been devalued and stolen by the voracious form of capitalism that drives the internet, and which the internet drives in turn. There is less time these days for anything other than economic survival. The internet has moved seamlessly into the interstices of this situation, redistributing our minimum of free time into unsatisfying micro-installments, spread throughout the day. In the absence of time to physically and politically engage with our community the way many of us want to, the internet provides a cheap substitute: it gives us brief moments of pleasure and connection, tied up in the opportunity to constantly listen and speak. Under these circumstances, opinion stops being a first step toward something and starts seeming like an end in itself.

I started thinking about this when I was working as an editor at *Jezebel*, in 2014. I spent a lot of the day reading headlines on women's websites, most of which had by then adopted a feminist slant. In this realm, speech was constantly framed as a sort of intensely satisfying action: you'd get headlines like "Miley Cyrus Spoke Out About Gender Fluidity on Snapchat and It Was Everything" or "Amy Schumer's Speech About Body Confidence at the Women's Magazine Awards Ceremony Will Have You in Tears." Forming an opinion was also framed as a sort of action: blog posts offered people guidance on how to feel about online controversies or particular scenes on TV. Even identity itself seemed to take on these valences. Merely to exist as a feminist was to be doing some important work. These ideas have intensified and gotten more complicated in the Trump era, in which, on the one hand, people like me are busy expressing anguish online and mostly affecting *nothing* and on the other, more actual and rapid change has come

from the internet than ever before. In the turbulence that followed the Harvey Weinstein revelations, women's speech swayed public opinion and led directly to change. People with power were forced to reckon with their ethics; harassers and abusers were pushed out of their jobs. But even in this narrative, the importance of action was subtly elided. People wrote about women "speaking out" with prayerful reverence, as if speech itself could bring women freedom—as if better policies and economic redistribution and true investment from men weren't necessary, too.

Goffman observes the difference between doing something and *expressing* the doing of something, between feeling something and conveying a feeling. "The representation of an activity will vary in some degree from the activity itself and therefore inevitably misrepresent it," Goffman writes. (Take the experience of enjoying a sunset versus the experience of communicating to an audience that you're enjoying a sunset, for example.) The internet is engineered for this sort of misrepresentation; it's designed to encourage us to create certain impressions rather than allowing these impressions to arise "as an incidental by-product of [our] activity." This is why, with the internet, it's so easy to stop trying to be decent, or reasonable, or politically engaged—and start trying merely to *seem* so.

As the value of speech inflates even further in the online attention economy, this problem only gets worse. I don't know what to do with the fact that I myself continue to benefit from all this: that my career is possible in large part because of the way the internet collapses identity, opinion, and action—and that I, as a writer whose work is mostly critical and often written in first person, have some inherent stake in justifying the dubious practice of spending all day trying to figure out what you think. As a reader, of course, I'm grateful for people who help me understand things, and I'm glad that they—and I—can be paid to do so. I am glad, too, for the way the internet has given an audience to writers who previously might have been shut out of the industry,

or kept on its sidelines: I'm one of them. But you will never catch me arguing that professional opinion-havers in the age of the internet are, on the whole, a force for good.

In April 2017, the *Times* brought a millennial writer named Bari Weiss onto its opinion section as both a writer and an editor. Weiss had graduated from Columbia, and had worked as an editor at *Tablet* and then at *The Wall Street Journal*. She leaned conservative, with a Zionist streak. At Columbia, she had cofounded a group called Columbians for Academic Freedom, hoping to pressure the university into punishing a pro-Palestinian professor who had made her feel "intimidated," she told NPR in 2005.

At the *Times*, Weiss immediately began launching columns from a rhetorical and political standpoint of high-strung defensiveness, disguised with a veneer of levelheaded nonchalance. "Victimhood, in the intersectional way of seeing the world, is akin to sainthood; power and privilege are profane," she wrote—a bit of elegant phrasing in a piece that warned the public of the rampant anti-Semitism evinced, apparently, by a minor activist clusterfuck, in which the organizers of the Chicago Dyke March banned Star of David flags. She wrote a column slamming the organizers of the Women's March over a few social media posts expressing support for Assata Shakur and Louis Farrakhan. This, she argued, was troubling evidence that progressives, just like conservatives, were unable to police their internal hate. (Both-sides arguments like this are always appealing to people who wish to seem both contrarian and intellectually superior; this particular one required ignoring the fact that liberals remained obsessed with "civility" while the Republican president was actively endorsing violence at every turn. Later on, when *Tablet* published an investigation into the Women's March organizers who maintained disconcerting ties to the Nation of Islam, these organizers were criticized by liberals, who truly do not lack the self-policing in-

stinct; in large part because the left does take hate seriously, the Women's March effectively splintered into two groups.) Often, Weiss's columns featured aggrieved predictions of how her bold, independent thinking would make her opponents go crazy and attack her. "I will inevitably get called a racist," she proclaimed in one column, titled "Three Cheers for Cultural Appropriation." "I'll be accused of siding with the alt-right or tarred as Islamophobic," she wrote in another column. Well, sure.

Though Weiss often argued that people should get more comfortable with those who offended or disagreed with them, she seemed mostly unable to take her own advice. During the Winter Olympics in 2018, she watched the figure skater Mirai Nagasu land a triple axel—the first American woman to do so in Olympic competition—and tweeted, in a very funny attempt at a compliment, "Immigrants: they get the job done." Because Nagasu was actually born in California, Weiss was immediately shouted down. This is what happens online when you do something offensive: when I worked at *Jezebel*, people shouted me down on Twitter about five times a year over things I had written or edited, and sometimes outlets published pieces about our mistakes. This was often overwhelming and unpleasant, but it was always useful. Weiss, for her part, tweeted that the people calling her racist tweet racist were a "sign of civilization's end." A couple of weeks later, she wrote a column called "We're All Fascists Now," arguing that angry liberals were creating a "moral flattening of the earth." At times it seems that Weiss's main strategy is to make an argument that's bad enough to attract criticism, and then to cherry-pick the worst of that criticism into the foundation for another bad argument. Her worldview requires the specter of a vast, angry, inferior mob.

It's of course true that there are vast, angry mobs on the internet. Jon Ronson wrote the book *So You've Been Publicly Shamed* about this in 2015. "We became keenly watchful for transgressions," he writes, describing the state of Twitter around 2012.

“After a while it wasn’t just transgressions we were keenly watchful for. It was misspeakings. Fury at the terribleness of other people had started to consume us a lot. . . . In fact, it felt weird and empty when there *wasn’t* anyone to be furious about. The days between shamings felt like days picking at fingernails, treading water.” Web 2.0 had curdled; its organizing principle was shifting. The early internet had been constructed around lines of affinity, and whatever good spaces remain on the internet are still the product of affinity and openness. But when the internet moved to an organizing principle of *opposition*, much of what had formerly been surprising and rewarding and curious became tedious, noxious, and grim.

This shift partly reflects basic social physics. Having a mutual enemy is a quick way to make a friend—we learn this as early as elementary school—and politically, it’s much easier to organize people against something than it is to unite them in an affirmative vision. And, within the economy of attention, conflict always gets more people to look. Gawker Media thrived on antagonism: its flagship site made enemies of everyone; *Deadspin* targeted ESPN, *Jezebel* the world of women’s magazines. There was a brief wave of sunny, saccharine, profitable internet content—the OMG era of *BuzzFeed*, the rise of sites like *Upworthy*—but it ended in 2014 or so. Today, on Facebook, the most-viewed political pages succeed because of a commitment to constant, aggressive, often unhinged opposition. Beloved, oddly warmhearted websites like *The Awl*, *The Toast*, and *Grantland* have all been shuttered; each closing has been a reminder that an open-ended, affinity-based, generative online identity is hard to keep alive.

That opposition looms so large on the internet can be good and useful and even revolutionary. Because of the internet’s tilt toward decontextualization and frictionlessness, a person on social media can seem to matter as much as whatever he’s set himself against. Opponents can meet on suddenly (if temporarily) even ground. *Gawker* covered the accusations against Louis C.K.

and Bill Cosby years before the mainstream media would take sexual misconduct seriously. The Arab Spring, Black Lives Matter, and the movement against the Dakota Access Pipeline challenged and overturned long-standing hierarchies through the strategic deployment of social media. The Parkland teenagers were able to position themselves as opponents of the entire GOP.

But the appearance of a more level playing field is not the fact of it, and everything that happens on the internet bounces and refracts. At the same time that ideologies that lead toward equality and freedom have gained power through the internet’s open discourse, existing power structures have solidified through a vicious (and very online) opposition to this encroachment. In her 2017 book, *Kill All Normies*—a project of accounting for the “online battles that may otherwise be forgotten but have nevertheless shaped culture and ideas in a profound way”—the writer Angela Nagle argues that the alt-right coalesced in response to increasing cultural power on the left. Gamergate, she writes, brought together a “strange vanguard of teenage gamers, pseudonymous swastika-posting anime lovers, ironic *South Park* conservatives, anti-feminist pranksters, nerdish harassers and meme-making trolls” to form a united front against the “earnestness and moral self-flattery of what felt like a tired liberal intellectual conformity.” The obvious hole in the argument is the fact that what Nagle identifies as the center of this liberal conformity—college activist movements, obscure Tumblr accounts about mental health and arcane sexualities—are frequently derided by liberals, and have never been nearly as powerful as those who detest them would like to think. The Gamergaters’ worldview was not actually endangered; they just had to *believe* it was—or to pretend it was, and wait for a purportedly leftist writer to affirm them—in order to lash out and remind everyone what they could do.

Many Gamergaters cut their expressive teeth on 4chan, a message board that adopted as one of its mottos the phrase “There are no girls on the internet.” “This rule does not mean what you think

it means,” wrote one 4chan poster, who went, as most of them did, by the username Anonymous. “In real life, people like you for being a girl. They want to fuck you, so they pay attention to you and they pretend what you have to say is interesting, or that you are smart or clever. On the Internet, we don’t have the chance to fuck you. This means the advantage of being a ‘girl’ does not exist. You don’t get a bonus to conversation just because I’d like to put my cock in you.” He explained that women could get their unfair social advantage back by posting photos of their tits on the message board: “This is, and should be, degrading for you.”

Here was the opposition principle in action. Through identifying the effects of women’s systemic objectification as some sort of vagina-supremacist witchcraft, the men that congregated on 4chan gained an identity, and a useful common enemy. Many of these men *had*, likely, experienced consequences related to the “liberal intellectual conformity” that is popular feminism: as the sexual marketplace began to equalize, they suddenly found themselves unable to obtain sex by default. Rather than work toward other forms of self-actualization—or attempt to make themselves genuinely desirable, in the same way that women have been socialized to do at great expense and with great sincerity for all time—they established a group identity that centered on anti-woman virulence, on telling women who happened to stumble across 4chan that “the only interesting thing about you is your naked body. tl;dr: tits or GET THE FUCK OUT.”

In the same way that it behooved these trolls to credit women with a maximum of power that they did not actually possess, it sometimes behooved women, on the internet, to do the same when they spoke about trolls. At some points while I worked at *Jezebel*, it would have been easy to enter into one of these situations myself. Let’s say a bunch of trolls sent me threatening emails—an experience that wasn’t exactly common, as I have been “lucky,” but wasn’t rare enough to surprise me. The economy of online attention would suggest that I write a column about

those trolls, quote their emails, talk about how the experience of being threatened constitutes a definitive situation of being a woman in the world. (It would be acceptable for me to do this *even though* I have never been hacked or swatted or Gamergated, never had to move out of my house to a secure location, as so many other women have.) My column about trolling would, of course, attract an influx of trolling. Then, having proven my point, maybe I’d go on TV and talk about the situation, and then I would get trolled even more, and then I could go on defining myself in reference to trolls forever, positioning them as inexorable and monstrous, and they would return the favor in the interest of their own ideological advancement, and this whole situation could continue until we all died.

There is a version of this mutual escalation that applies to any belief system, which brings me back to Bari Weiss and all the other writers who have fashioned themselves as brave contrarians, building entire arguments on random protests and harsh tweets, making themselves deeply dependent on the people who hate them, the people they hate. It’s ridiculous, and at the same time, here I am writing this essay, doing the same thing. It is nearly impossible, today, to separate engagement from magnification. (Even declining to engage can turn into magnification: when people targeted in Pizzagate as Satanist pedophiles took their social media accounts private, the Pizzagaters took this as proof that they had been right.) Trolls and bad writers and the president know better than anyone: when you call someone terrible, you just end up promoting their work.

The political philosopher Sally Scholz separates solidarity into three categories. There’s social solidarity, which is based on common experience; civic solidarity, which is based on moral obligation to a community; and political solidarity, which is based on a shared commitment to a cause. These forms of solidarity overlap,

but they're distinct from one another. What's political, in other words, doesn't also have to be personal, at least not in the sense of firsthand experience. You don't need to step in shit to understand what stepping in shit feels like. You don't need to have directly suffered at the hands of some injustice in order to be invested in bringing that injustice to an end.

But the internet brings the "I" into everything. The internet can make it seem that supporting someone means literally sharing in their experience—that solidarity is a matter of *identity* rather than politics or morality, and that it's best established at a point of maximum mutual vulnerability in everyday life. Under these terms, instead of expressing morally obvious solidarity with the struggle of black Americans under the police state or the plight of fat women who must roam the earth to purchase stylish and thoughtful clothing, the internet would encourage me to express solidarity through inserting my own identity. *Of course* I support the black struggle because *I*, myself, as a woman of Asian heritage, have *personally* been injured by white supremacy. (In fact, as an Asian woman, part of a minority group often deemed white-adjacent, I have benefited from American anti-blackness on just as many occasions.) *Of course* I understand the difficulty of shopping as a woman who is overlooked by the fashion industry because *I*, myself, have *also* somehow been marginalized by this industry. This framework, which centers the self in an expression of support for others, is not ideal.

The phenomenon in which people take more comfort in a sense of injury than a sense of freedom governs many situations where people are objectively *not* being victimized on a systematic basis. For example, men's rights activists have developed a sense of solidarity around the absurd claim that men are second-class citizens. White nationalists have brought white people together through the idea that white people are endangered, specifically white men—this at a time when 91 percent of Fortune 500 CEOs are white men, when white people make up 90 percent of elected

American officials and an overwhelming majority of top decision-makers in music, publishing, television, movies, and sports.

Conversely, and crucially, the dynamic also applies in situations where claims of vulnerability are legitimate and historically entrenched. The greatest moments of feminist solidarity in recent years have stemmed not from an affirmative vision but from articulating extreme versions of the low common denominator of male slight. These moments have been world-altering: #YesAllWomen, in 2014, was the response to Elliot Rodger's Isla Vista massacre, in which he killed six people and wounded fourteen in an attempt to exact revenge on women for rejecting him. Women responded to this story with a sense of nauseating recognition: mass violence is nearly always linked to violence toward women, and for women it is something approaching a universal experience to have placated a man out of the real fear that he will hurt you. In turn, some men responded with the entirely unnecessary reminder that "not all men" are like that. (I was once hit with "not all men" right after a stranger yelled something obscene at me; the guy I was with noted my displeasure and helpfully reminded me that not all men are jerks.) Women began posting stories on Twitter and Facebook with #YesAllWomen to make an obvious but important point: not all men have made women fearful, but yes, all women have experienced fear because of men. #MeToo, in 2017, came in the weeks following the Harvey Weinstein revelations, as the floodgates opened and story after story after story rolled out about the subjugation women had experienced at the hands of powerful men. Against the normal forms of disbelief and rejection these stories meet with—it can't possibly be *that* bad; something about *her* telling *that* story seems suspicious—women anchored one another, establishing the breadth and inescapability of male abuse of power through speaking simultaneously and adding #MeToo.

In these cases, multiple types of solidarity seemed to naturally meld together. It was women's individual experiences of victim-

ization that produced our widespread moral and political opposition to it. And at the same time, there was something about the hashtag itself—its design, and the ways of thinking that it affirms and solidifies—that both erased the variety of women’s experiences and made it seem as if the crux of feminism was this articulation of vulnerability itself. A hashtag is specifically designed to remove a statement from context and to position it as part of an enormous singular thought. A woman participating in one of these hashtags becomes visible at an inherently predictable moment of male aggression: the time her boss jumped her, or the night a stranger followed her home. The rest of her life, which is usually far less predictable, remains unseen. Even as women have attempted to use #YesAllWomen and #MeToo to regain control of a narrative, these hashtags have at least partially reified the thing they’re trying to eradicate: the way that womanhood can feel like a story of loss of control. They have made feminist solidarity and shared vulnerability seem inextricable, as if we were incapable of building solidarity around anything else. What we have in common is obviously essential, but it’s the differences between women’s stories—the factors that allow some to survive, and force others under—that illuminate the vectors that lead to a better world. And, because there is no room or requirement in a tweet to add a disclaimer about individual experience, and because hashtags subtly equate disconnected statements in a way that can’t be controlled by those speaking, it has been even easier for #MeToo critics to claim that women must themselves think that going on a bad date is the same as being violently raped.

What’s amazing is that things like hashtag design—these essentially ad hoc experiments in digital architecture—have shaped so much of our political discourse. Our world would be different if Anonymous hadn’t been the default username on 4chan, or if every social media platform didn’t center on the personal profile, or if YouTube algorithms didn’t show viewers increasingly extreme content to retain their attention, or if hashtags and retweets

simply didn’t exist. It’s because of the hashtag, the retweet, and the profile that solidarity on the internet gets inextricably tangled up with visibility, identity, and self-promotion. It’s telling that the most mainstream gestures of solidarity are pure representation, like viral reposts or avatar photos with cause-related filters, and meanwhile the *actual* mechanisms through which political solidarity is enacted, like strikes and boycotts, still exist on the fringe. The extremes of performative solidarity are all transparently embarrassing: a Christian internet personality urging other conservatives to tell Starbucks baristas that their name is “Merry Christmas,” or Nev Schulman from the TV show *Catfish* taking a selfie with a hand over his heart in an elevator and captioning it “A real man shows his strength through patience and honor. This elevator is abuse free.” (Schulman punched a girl in college.) The demonstrative celebration of black women on social media—white people tweeting “black women will save America” after elections, or Mark Ruffalo tweeting that he said a prayer and God answered as a black woman—often hints at a bizarre need on the part of white people to personally participate in an ideology of equality that ostensibly requires them to chill out. At one point in *The Presentation of Self*, Goffman writes that the audience’s way of shaping a role for the performer can become more elaborate than the performance itself. This is what the online expression of solidarity sometimes feels like—a manner of listening so extreme and performative that it often turns into the show.

The final, and possibly most psychologically destructive, distortion of the social internet is its distortion of scale. This is not an accident but an essential design feature: social media was constructed around the idea that a thing is important insofar as it is important to you. In an early internal memo about the creation of Facebook’s News Feed, Mark Zuckerberg observed, already beyond parody, “A squirrel dying in front of your house may be more

relevant to your interests right now than people dying in Africa.” The idea was that social media would give us a fine-tuned sort of control over what we looked at. What resulted was a situation where we—first as individuals, and then inevitably as a collective—are essentially unable to exercise control at all. Facebook’s goal of showing people only what they were interested in seeing resulted, within a decade, in the effective end of shared civic reality. And this choice, combined with the company’s financial incentive to continually trigger heightened emotional responses in its users, ultimately solidified the current norm in news media consumption: today we mostly consume news that corresponds with our ideological alignment, which has been fine-tuned to make us feel self-righteous and also mad.

In *The Attention Merchants*, Tim Wu observes that technologies designed to increase control over our attention often have the opposite effect. He uses the TV remote control as one example. It made flipping through channels “practically nonvolitional,” he writes, and put viewers in a “mental state not unlike that of a newborn or a reptile.” On the internet, this dynamic has been automated and generalized in the form of endlessly varied but somehow monotonous social media feeds—these addictive, numbing fire hoses of information that we aim at our brains for much of the day. In front of the timeline, as many critics have noted, we exhibit classic reward-seeking lab-rat behavior, the sort that’s observed when lab rats are put in front of an unpredictable food dispenser. Rats will eventually stop pressing the lever if their device dispenses food regularly or not at all. But if the lever’s rewards are rare and irregular, the rats will never stop pressing it. In other words, it is *essential* that social media is mostly unsatisfying. That is what keeps us scrolling, scrolling, pressing our lever over and over in the hopes of getting some fleeting sensation—some momentary rush of recognition, flattery, or rage.

Like many among us, I have become acutely conscious of the way my brain degrades when I strap it in to receive the full barrage

of the internet—these unlimited channels, all constantly reloading with new information: births, deaths, boasts, bombings, jokes, job announcements, ads, warnings, complaints, confessions, and political disasters blitzing our frayed neurons in huge waves of information that pummel us and then are instantly replaced. This is an awful way to live, and it is wearing us down quickly. At the end of 2016, I wrote a blog post for *The New Yorker* about the cries of “worst year ever” that were then flooding the internet. There had been terrorist attacks all over the world, and the Pulse shooting in Orlando. David Bowie, Prince, and Muhammad Ali had died. More black men had been executed by police who could not control their racist fear and hatred: Alton Sterling was killed in the Baton Rouge parking lot where he was selling CDs; Philando Castile was murdered as he reached for his legal-carry permit during a routine traffic stop. Five police officers were killed in Dallas at a protest against this police violence. Donald Trump was elected president of the United States. The North Pole was thirty-six degrees hotter than normal. Venezuela was collapsing; families starved in Yemen. In Aleppo, a seven-year-old girl named Bana Alabed was tweeting her fears of imminent death. And in front of this backdrop, there were all of *us*—our stupid selves, with our stupid frustrations, our lost baggage and delayed trains. It seemed to me that this sense of punishing oversaturation would persist no matter what was in the news. There was no limit to the amount of misfortune a person could take in via the internet, I wrote, and there was no way to calibrate this information correctly—no guidebook for how to expand our hearts to accommodate these simultaneous scales of human experience, no way to teach ourselves to separate the banal from the profound. The internet was dramatically increasing our ability to know about things, while our ability to *change* things stayed the same, or possibly shrank right in front of us. I had started to feel that the internet would only ever induce this cycle of heartbreak and hardening—a hyper-engagement that would make less sense every day.

But the worse the internet gets, the more we appear to crave it—the more it gains the power to shape our instincts and desires. To guard against this, I give myself arbitrary boundaries—no Instagram stories, no app notifications—and rely on apps that shut down my Twitter and Instagram accounts after forty-five minutes of daily use. And still, on occasion, I'll disable my social media blockers, and I'll sit there like a rat pressing the lever, like a woman repeatedly hitting myself on the forehead with a hammer, masturbating through the nightmare until I finally catch the gasoline whiff of a good meme. The internet is still so young that it's easy to retain some subconscious hope that it all might still add up to something. We remember that at one point this all felt like butterflies and puddles and blossoms, and we sit patiently in our festering inferno, waiting for the internet to turn around and surprise us and get good again. But it won't. The internet is governed by incentives that make it impossible to be a full person while interacting with it. In the future, we will inevitably be cheapened. Less and less of us will be left, not just as individuals but also as community members, as a collective of people facing various catastrophes. Distraction is a "life-and-death matter," Jenny Odell writes in *How to Do Nothing*. "A social body that can't concentrate or communicate with itself is like a person who can't think and act."

Of course, people have been carping in this way for many centuries. Socrates feared that the act of writing would "create forgetfulness in the learners' souls." The sixteenth-century scientist Conrad Gessner worried that the printing press would facilitate an "always on" environment. In the eighteenth century, men complained that newspapers would be intellectually and morally isolating, and that the rise of the novel would make it difficult for people—specifically women—to differentiate between fiction and fact. We worried that radio would drive children to distraction, and later that TV would erode the careful attention required by radio. In 1985, Neil Postman observed that the

American desire for constant entertainment had become toxic, that television had ushered in a "vast descent into triviality." The difference is that, today, there is nowhere further to go. Capitalism has no land left to cultivate but the self. Everything is being cannibalized—not just goods and labor, but personality and relationships and attention. The next step is complete identification with the online marketplace, physical and spiritual inseparability from the internet: a nightmare that is already banging down the door.

What could put an end to the worst of the internet? Social and economic collapse would do it, or perhaps a series of antitrust cases followed by a package of hard regulatory legislation that would somehow also dismantle the internet's fundamental profit model. At this point it's clear that collapse will almost definitely come first. Barring that, we've got nothing except our small attempts to retain our humanity, to act on a model of actual selfhood, one that embraces culpability, inconsistency, and insignificance. We would have to think very carefully about what we're getting from the internet, and how much we're giving it in return. We'd have to care less about our identities, to be deeply skeptical of our own unbearable opinions, to be careful about when opposition serves us, to be properly ashamed when we can't express solidarity without putting ourselves first. The alternative is unspeakable. But you know that—it's already here.