

# Facebook Is a Doomsday Machine

*Adrienne LaFrance*

The Doomsday Machine was never supposed to exist. It was meant to be a thought experiment that went like this: Imagine a device built with the sole purpose of destroying all human life. Now suppose that machine is buried deep underground, but connected to a computer, which is in turn hooked up to sensors in cities and towns across the United States.

The sensors are designed to sniff out signs of the impending apocalypse—not to prevent the end of the world, but to complete it. If radiation levels suggest nuclear explosions in, say, three American cities simultaneously, the sensors notify the Doomsday Machine, which is programmed to detonate several nuclear warheads in response. At that point, there is no going back. The fission chain reaction that produces an atomic explosion is initiated enough times over to extinguish all life on Earth. There is a terrible flash of light, a great booming sound, then a sustained roar. We have a word for the scale of destruction that the Doomsday Machine would unleash: *megadeath*.

Nobody is pining for megadeath. But megadeath is not the only thing that makes the Doomsday Machine petrifying. The real terror is in its autonomy, this idea that it would be programmed to detect a series of environmental inputs, then to act, without human interference. “There is no chance of human intervention, control, and final decision,” wrote the military strategist Herman Kahn in his 1960 book, *On Thermonuclear War*, which laid out the hypothetical for a Doomsday Machine. The concept was to render nuclear war unwinnable, and therefore unthinkable.

Kahn concluded that automating the extinction of all life on Earth would be immoral. Even an infinitesimal risk of error is too great to justify the Doomsday Machine’s existence. “And even if we give up the computer and make the Doomsday Machine reliably controllable by decision makers,” Kahn wrote, “it is still not controllable enough.” No machine should be that powerful by itself—but no one person should be either.

The Soviets really did make a version of the Doomsday Machine during the Cold War. They nicknamed it “Dead Hand.” But so far, somewhat miraculously, we have figured out how to live with the bomb. Now we need to learn how to survive the social web.

People tend to complain about Facebook as if something recently curdled. There’s a notion that the social web was once useful, or at least that it could have been good, if only we had pulled a few levers: some moderation and fact-checking here, a bit of regulation there, perhaps a federal antitrust lawsuit. But that’s far too sunny and

shortsighted a view. Today's social networks, Facebook chief among them, were built to encourage the things that make them so harmful. It is in their very architecture.

I've been thinking for years about what it would take to make the social web magical in all the right ways—less extreme, less toxic, more true—and I realized only recently that I've been thinking far too narrowly about the problem. I've long wanted Mark Zuckerberg to admit that [Facebook is a media company](#), to take responsibility for the informational environment he created in the same way that the editor of a magazine would. (I pressed him on this [once](#) and he laughed.) In recent years, as Facebook's mistakes have compounded and its reputation has tanked, it has become clear that negligence is only part of the problem. No one, not even Mark Zuckerberg, can control the product he made. I've come to realize that Facebook is not a media company. It's a Doomsday Machine.

[Read: Breaking up Facebook isn't enough](#)

The social web is doing exactly what it was built for. Facebook does not exist to seek truth and report it, or to improve civic health, or to hold the powerful to account, or to represent the interests of its users, though these phenomena may be occasional by-products of its existence. The company's early mission was to "give people the power to share and make the world more open and connected." Instead, it took the concept of "community" and sapped it of all moral meaning. [The rise of QAnon](#), for example, is one of the social web's logical conclusions. That's because Facebook—along with Google and YouTube—is perfect for amplifying and spreading disinformation at lightning speed to global audiences. Facebook is an agent of government propaganda, targeted harassment, terrorist recruitment, emotional manipulation, and genocide—a world-historic weapon that lives not underground, but in a Disneyland-inspired campus in Menlo Park, California.

The giants of the social web—Facebook and its subsidiary Instagram; Google and its subsidiary YouTube; and, to a lesser extent, Twitter—have achieved success by being dogmatically value-neutral in their pursuit of what I'll call *megascale*. Somewhere along the way, Facebook decided that it needed not just a very large user base, but a tremendous one, unprecedented in size. That decision set Facebook on a path to escape velocity, to a tipping point where it can harm society just by existing.

Limitations to the Doomsday Machine comparison are obvious: Facebook cannot in an instant reduce a city to ruins the way a nuclear bomb can. And whereas the Doomsday Machine was conceived of as a world-ending device so as to forestall the end of the world; Facebook started because a semi-inebriated Harvard undergrad was bored one night. But the stakes are still life-and-death. Megascale is nearly the existential threat that megadeath is. No single machine should be able to control the fate of the world's population—and that's what both the Doomsday Machine and Facebook are built to do.

The cycle of harm perpetuated by Facebook's scale-at-any-cost business model is plain to see. Scale and engagement are valuable to Facebook because they're

valuable to advertisers. These incentives lead to design choices such as reaction buttons that encourage users to engage easily and often, which in turn encourage users to share ideas that will provoke a strong response. Every time you click a reaction button on Facebook, an algorithm records it, and sharpens its portrait of who you are. The hyper-targeting of users, made possible by reams of their personal data, creates the perfect environment for manipulation—by advertisers, by political campaigns, by emissaries of disinformation, and of course by Facebook itself, which ultimately controls what you see and what you don't see on the site. Facebook has enlisted a corps of approximately 15,000 moderators, people [paid to watch unspeakable things](#)—murder, gang rape, and other depictions of graphic violence that wind up on the platform. Even as Facebook has insisted that it is a value-neutral vessel for the material its users choose to publish, moderation is a lever the company has tried to pull again and again. But there aren't enough moderators speaking enough languages, working enough hours, to stop the biblical flood of shit that Facebook unleashes on the world, because 10 times out of 10, the algorithm is faster and more powerful than a person. At megascale, this algorithmically warped personalized informational environment is extraordinarily difficult to moderate in a meaningful way, and extraordinarily dangerous as a result.

These dangers are not theoretical, and they're exacerbated by megascale, which makes the platform a tantalizing place to experiment on people. Facebook has conducted social-contagion [experiments](#) on its users without telling them. Facebook has acted as [a force for digital colonialism](#), attempting to become the de facto (and only) experience of the internet for people all over the world. Facebook [has bragged](#) about its ability to influence the outcome of elections. Unlawful militant groups [use Facebook to organize](#). Government officials use Facebook [to mislead](#) their own citizens, and to tamper with elections. Military officials have [exploited](#) Facebook's complacency to carry out genocide. Facebook inadvertently [auto-generated](#) jaunty recruitment videos for the Islamic State featuring anti-Semitic messages and burning American flags.

[Read: Mark Zuckerberg doesn't understand journalism](#)

Even after U.S. intelligence agencies [identified Facebook](#) as a main battleground for information warfare and foreign interference in the 2016 election, the company has failed to stop the spread of extremism, hate speech, propaganda, disinformation, and conspiracy theories on its site. Neo-Nazis [stayed active on Facebook](#) by taking out ads even after they were formally banned. And it wasn't until October of this year, for instance, that Facebook announced it would remove groups, pages, and Instagram accounts devoted to QAnon, as well as any posts denying the Holocaust. (Previously Zuckerberg had defended Facebook's decision not to remove disinformation about the Holocaust, saying of Holocaust deniers, "I don't think that they're *intentionally* getting it wrong." He [later clarified](#) that he didn't mean to defend Holocaust deniers.) Even so, Facebook routinely sends emails to users recommending the newest QAnon groups. White supremacists and deplatformed MAGA trolls may flock to smaller social platforms such as Gab and Parler, but these platforms offer little aside from a narrative of martyrdom without megascale.

In the days after the 2020 presidential election, Zuckerberg authorized a tweak to the Facebook algorithm so that high-accuracy news sources such as NPR would receive preferential visibility in people's feeds, and hyper-partisan pages such as *Breitbart News's* and Occupy Democrats' would be buried, [according to \*The New York Times\*](#), offering proof that Facebook could, if it wanted to, turn a dial to reduce disinformation—and offering a reminder that Facebook has the power to flip a switch and change what billions of people see online.

The decision to touch the dial was highly unusual for Facebook. Think about it this way: The Doomsday Machine's sensors detected something harmful in the environment and chose not to let its algorithms automatically blow it up across the web as usual. This time a human intervened to mitigate harm. The only problem is that reducing the prevalence of content that Facebook calls “bad for the world” also reduces people's engagement with the site. In its experiments with human intervention, the *Times* reported, Facebook calibrated the dial so that *just enough* harmful content stayed in users' news feeds to keep them coming back for more.

Facebook's stated mission—to make the world more open and connected—has always seemed, to me, phony at best, and imperialist at worst. After all, today's empires are born on the web. Facebook is a borderless nation-state, with a population of users nearly as big as China and India combined, and it is governed largely by secret algorithms. Hillary Clinton [told me](#) earlier this year that talking to Zuckerberg feels like negotiating with the authoritarian head of a foreign state. “This is a global company that has huge influence in ways that we're only beginning to understand,” she said.

I recalled Clinton's warning a few weeks ago, when Zuckerberg defended the decision not to suspend Steve Bannon from Facebook after he argued, in essence, for the beheading of two senior U.S. officials, the infectious-disease doctor Anthony Fauci and FBI Director Christopher Wray. The episode got me thinking about a question that's unanswerable but that I keep asking people anyway: How much real-world violence would never have happened if Facebook didn't exist? One of the people I've asked is Joshua Geltzer, a former White House counterterrorism official who is now teaching at Georgetown Law. In counterterrorism circles, he told me, people are fond of pointing out how good the United States has been at keeping terrorists out since 9/11. That's wrong, he said. In fact, “terrorists are entering every single day, every single hour, every single minute” through Facebook.

The website that's perhaps best known for encouraging mass violence is the image board 4chan—which was followed by 8chan, which then became 8kun. These boards are infamous for being the sites where multiple mass-shooting suspects have shared manifestos before homicide sprees. The few people who are willing to defend these sites unconditionally do so from a position of free-speech absolutism. That argument is worthy of consideration. But there's something architectural about the site that merits attention, too: There are no algorithms on 8kun, only a community of users who post what they want. People use 8kun to publish abhorrent ideas, but at least the community isn't pretending to be something it's not. The biggest social platforms claim to be similarly neutral and pro-free speech when in fact no two people see the

same feed. Algorithmically tweaked environments feed on user data and manipulate user experience, and not ultimately for the purpose of serving the user. Evidence of real-world violence can be easily traced back to both Facebook and 8kun. But 8kun doesn't manipulate its users or the informational environment they're in. Both sites are harmful. But Facebook might actually be worse for humanity.

[Read: How Facebook works for Trump](#)

“What a dreadful set of choices when you frame it that way,” Geltzer told me when I put this question to him in another conversation. “The idea of a free-for-all sounds really bad until you see what the purportedly moderated and curated set of platforms is yielding ... It may not be blood onscreen, but it can really do a lot of damage.”

In previous eras, U.S. officials could at least study, say, Nazi propaganda during World War II, and fully grasp what the Nazis wanted people to believe. Today, “it’s not a filter bubble; it’s a filter shroud,” Geltzer said. “I don’t even know what others with personalized experiences are seeing.” Another expert in this realm, Mary McCord, the legal director at the Institute for Constitutional Advocacy and Protection at Georgetown Law, told me that she thinks 8kun may be more blatant in terms of promoting violence but that Facebook is “in some ways way worse” because of its reach. “There’s no barrier to entry with Facebook,” she said. “In every situation of extremist violence we’ve looked into, we’ve found Facebook postings. And that reaches *tons* of people. The broad reach is what brings people into the fold and normalizes extremism and makes it mainstream.” In other words, it’s the megascale that makes Facebook so dangerous.

Looking back, it can seem like Zuckerberg’s path to world domination was inevitable. There’s the computerized version of Risk he coded in ninth grade; his long-standing interest in the Roman empire; his obsession with information flow and human psychology. There’s the story of his first bona fide internet scandal, when he hacked into Harvard’s directory and lifted photos of students without their permission to make the hot-or-not-style website FaceMash. (“Child’s play” was how Zuckerberg [later described](#) the ease with which he broke into Harvard’s system.) There’s the disconnect between his lip service to privacy and the way Facebook actually works. (Here’s Zuckerberg in a private chat with a friend years ago, on the mountain of data he’d obtained from Facebook’s early users: “I have over 4,000 emails, pictures, addresses ... People just submitted it. I don’t know why. They ‘trust me.’ Dumb fucks.”) At various points over the years, he’s listed the following interests in his Facebook profile: Eliminating Desire, Minimalism, Making Things, Breaking Things, Revolutions, Openness, Exponential Growth, Social Dynamics, Domination.

Facebook’s megascale gives Zuckerberg an unprecedented degree of influence over the global population. If he isn’t the most powerful person on the planet, he’s very near the top. “It’s insane to have that much speechifying, silencing, and permitting power, not to mention being the ultimate holder of algorithms that determine the virality of anything on the internet,” Geltzer told me. “The thing he oversees has such an effect on cognition and people’s beliefs, which can change what they do with

their nuclear weapons or their dollars.”

Facebook’s [new oversight board](#), formed in response to backlash against the platform and tasked with making decisions concerning moderation and free expression, is an extension of that power. “The first 10 decisions they make will have more effect on speech in the country and the world than the next 10 decisions rendered by the U.S. Supreme Court,” Geltzer said. “That’s power. That’s real power.”

In 2005, the year I joined Facebook, the site still billed itself as an online directory to “Look up people at your school. See how people know each other. Find people in your classes and groups.” That summer, in Palo Alto, Zuckerberg gave an [interview](#) to a young filmmaker, who later posted the clip to YouTube. In it, you can see Zuckerberg still figuring out what Facebook is destined to be. The conversation is a reminder of the improbability of Zuckerberg’s youth when he launched Facebook. (It starts with him asking, “Should I put the beer down?” He’s holding a red Solo cup.) Yet, at 21 years old, Zuckerberg articulated something about his company that has held true, to dangerous effect: Facebook is not a single place on the web, but rather, “a lot of different individual communities.”

Today that includes QAnon and other extremist groups. Back then, it meant mostly juvenile expressions of identity in groups such as “I Went to a Public School ... Bitch” and, at Harvard, referencing the neoclassical main library, “The We Need to Have Sex in Widener Before We Graduate Interest Group.” In that 2005 interview, Zuckerberg is asked about the future of Facebook, and his response feels, in retrospect, like a tragedy: “I mean, there doesn’t necessarily have to be more. Like, a lot of people are focused on taking over the world, or doing the biggest thing, getting the most users. I think, like, part of making a difference and doing something cool is focusing intensely ... I mean, I really just want to see everyone focus on college and create a really cool college-directory product that just, like, is very relevant for students and has a lot of information that people care about when they’re in college.”

[Read: What we wrote about Facebook 12 years ago](#)

The funny thing is: This localized approach is part of what made megascale possible. Early constraints around membership—the requirement at first that users attended Harvard, and then that they attended any Ivy League school, and then that they had an email address ending in .edu—offered a sense of cohesiveness and community. It made people feel more comfortable sharing more of themselves. And more sharing among clearly defined demographics was good for business. In 2004, Zuckerberg said Facebook ran advertisements only to cover server costs. But over the next two years Facebook completely upended and redefined the entire advertising industry. The pre-social web destroyed classified ads, but the one-two punch of Facebook and Google decimated local news and most of the magazine industry—publications fought in earnest for digital pennies, which had replaced print dollars, and social giants scooped them all up anyway. No news organization can compete with the megascale of the social web. It’s just too massive.

The on-again, off-again Facebook executive Chris Cox once talked about the “magic

number” for start-ups, and how after a company surpasses 150 employees, things go sideways. “I’ve talked to so many start-up CEOs that after they pass this number, weird stuff starts to happen,” he [said](#) at a conference in 2016. This idea comes from the anthropologist Robin Dunbar, who argued that 148 is the maximum number of stable social connections a person can maintain. If we were to apply that same logic to the stability of a social platform, what number would we find?

“I think the sweet spot is 20 to 20,000 people,” the writer and internet scholar Ethan Zuckerman, who has spent much of his adult life thinking about [how to build a better web](#), told me. “It’s hard to have any degree of real connectivity after that.”

In other words, if the Dunbar number for running a company or maintaining a cohesive social life is 150 people; the magic number for a functional social platform is maybe 20,000 people. Facebook now has [2.7 billion](#) monthly users.

On the precipice of Facebook’s exponential growth, in 2007, Zuckerberg said something in an interview with the *Los Angeles Times* that now takes on a much darker meaning: “The things that are most powerful aren’t the things that people would have done otherwise if they didn’t do them on Facebook. Instead, it’s the things that would never have happened otherwise.”

Of the many things humans are consistently terrible at doing, seeing the future is somewhere near the top of the list. This flaw became a preoccupation among Megadeath Intellectuals such as Herman Kahn and his fellow economists, mathematicians, and former military officers at the Rand Corporation in the 1960s.

Kahn and his colleagues helped invent modern futurism, which was born of the existential dread that the bomb ushered in, and hardened by the understanding that most innovation is horizontal in nature—a copy of what already exists, rather than wholly new. Real invention is extraordinarily rare, and far more disruptive.

The logician and philosopher Olaf Helmer-Hirschberg, who overlapped with Kahn at Rand and would later co-found the Institute for the Future, arrived in California after having fled the Nazis, an experience that gave his desire to peer into the future a particular kind of urgency. He argued that the acceleration of technological change had established the need for a new epistemological approach to fields such as engineering, medicine, the social sciences, and so on. “No longer does it take generations for a new pattern of living conditions to evolve,” he [wrote](#), “but we are going through several major adjustments in our lives, and our children will have to adopt continual adaptation as a way of life.” In 1965, he wrote a book called *Social Technology* that aimed to create a scientific methodology for predicting the future.

[Read: The silence of the never Facebookers](#)

In those same years, Kahn was dreaming up his own hypothetical machine to provide a philosophical framework for the new threats humanity faced. He called it the Doomsday Machine, and also the Doomsday-in-a-Hurry Machine, and also the Homicide Pact Machine. Stanley Kubrick famously borrowed the concept for the

1964 film *Dr. Strangelove*, the cinematic apotheosis of the fatalism that came with living on hair-trigger alert for nuclear annihilation.

Today's fatalism about the brokenness of the internet feels similar. We're still in the infancy of this century's triple digital revolution of the internet, smartphones, and the social web, and we find ourselves in a dangerous and unstable informational environment, powerless to resist forces of manipulation and exploitation that we know are exerted on us but remain mostly invisible. The Doomsday Machine offers a lesson: We should not accept this current arrangement. No single machine should be able to control so many people.

If the age of reason was, in part, a reaction to [the existence of the printing press](#), and 1960s futurism was a reaction to the atomic bomb, we need a new philosophical and moral framework for living with the social web—a new Enlightenment for the information age, and one that will carry us back to shared reality and empiricism.

Andrew Bosworth, one of Facebook's longtime executives, has compared Facebook to sugar—in that it is “delicious” but best enjoyed in moderation. In a memo originally posted to Facebook's internal network last year, he argued for a philosophy of personal responsibility. “My grandfather took such a stance towards bacon and I admired him for it,” Bosworth wrote. “And social media is likely much less fatal than bacon.” But viewing Facebook merely as a vehicle for individual consumption ignores the fact of what it is—a network. Facebook is also a business, and a place where people spend time with one another. Put it this way: If you owned a store and someone walked in and started shouting Nazi propaganda or recruiting terrorists near the cash register, would you, as the shop owner, tell all of the other customers you couldn't possibly intervene?

Anyone who is serious about mitigating the damage done to humankind by the social web should, of course, consider quitting Facebook and Instagram and Twitter and any other algorithmically distorted informational environments that manipulate people. But we need to adopt a broader view of what it will take to fix the brokenness of the social web. That will require challenging the logic of today's platforms—and first and foremost challenging the very concept of megascale as a way that humans gather. If megascale is what gives Facebook its power, and what makes it dangerous, collective action against the web as it is today is necessary for change. The web's existing logic tells us that social platforms are free in exchange for a feast of user data; that major networks are necessarily global and centralized; that moderators make the rules. None of that need be the case. We need people who dismantle these notions by building alternatives. And we need enough people to care about these other alternatives to break the spell of venture capital and mass attention that fuels megascale and creates fatalism about the web as it is now.

I still believe the internet is good for humanity, but that's despite the social web, not because of it. We must also find ways to repair the aspects of our society and culture that the social web has badly damaged. This will require intellectual independence, respectful debate, and the same rebellious streak that helped establish Enlightenment values centuries ago.

We may not be able to predict the future, but we do know how it is made: through flashes of rare and genuine invention, sustained by people's time and attention. Right now, too many people are allowing algorithms and tech giants to manipulate them, and reality is slipping from our grasp as a result. This century's Domsday Machine is here, and humming along.

It does not have to be this way.

[Adrienne LaFrance](#) is the executive editor of *The Atlantic*. She was previously a senior editor and staff writer at *The Atlantic*, and the editor of *TheAtlantic.com*.

Enjoy unlimited access to *The Atlantic* for less than \$1 per week. Thank you for reading *The Atlantic*.

[Sign in](#)