


How Much of Google And Facebook Is Fake? Turns Out, Most of It, Actually.

By [Max Read](#)



[@max_read](#)

Photo: Artwork by Ayatgali Tuleubek

In late November, the Justice Department unsealed indictments against eight people accused of fleecing advertisers of \$36 million in two of the largest digital ad-fraud operations ever uncovered. Digital advertisers tend to want two things: people to look at their ads and “premium” websites — i.e., established and legitimate publications — on which to host them.

The two schemes at issue in the case, dubbed Methbot and 3ve by the security researchers who found them, faked both. Hucksters infected 1.7 million computers with malware that remotely directed traffic to “spoofed” websites — [“empty websites designed for bot traffic”](#) that served up a video ad purchased from one of the internet’s vast programmatic ad-exchanges, but that were designed, according to the indictments, “to fool advertisers into thinking that an impression of their ad was served on a premium publisher site,” like that of *Vogue* or *The Economist*. Views, meanwhile, were faked by malware-infected computers with marvelously sophisticated techniques to imitate humans: bots [“faked clicks, mouse movements, and social network login information to masquerade as engaged human consumers.”](#) Some were sent to browse the internet to gather tracking cookies from other websites, just as a human visitor would have done through regular behavior. Fake people with fake cookies and fake social-media accounts, fake-moving their fake cursors, fake-clicking on fake websites — the fraudsters had essentially created a simulacrum of the internet, where the only real things were the ads.

How much of the internet is fake? Studies generally suggest that, year after year, less than 60 percent of web traffic is human; some years, according to some researchers, a healthy majority of it is bot. For a period of time in 2013, the *Times* [reported](#) this

year, a full half of YouTube traffic was “bots masquerading as people,” a portion so high that employees feared an inflection point after which YouTube’s systems for detecting fraudulent traffic would begin to regard bot traffic as real and human traffic as fake. They called this hypothetical event “the Inversion.”

In the future, when I look back from the high-tech gamer jail in which [President PewDiePie will have imprisoned me](#), I will remember 2018 as the year the internet passed the Inversion, not in some strict numerical sense, since bots already outnumber humans online more years than not, but in the perceptual sense. The internet has always played host in its dark corners to schools of catfish and embassies of Nigerian princes, but that darkness now pervades its every aspect: Everything that once seemed definitively and unquestionably real now seems slightly fake; everything that once seemed slightly fake now has the power and presence of the real. The “fakeness” of the post-Inversion internet is less a calculable falsehood and more a particular quality of experience — the uncanny sense that what you encounter online is not “real” but is also undeniably not “fake,” and indeed may be both at once, or in succession, as you turn it over in your head.

The metrics are fake.

Take something as seemingly simple as how we measure web traffic. Metrics should be the most real thing on the internet: They are countable, trackable, and verifiable, and their existence undergirds the advertising business that drives our biggest social and search platforms. Yet [not even Facebook](#), the world's greatest data-gathering organization, seems able to produce genuine figures. In October, [small advertisers filed suit against the social-media giant](#), accusing it of covering up, for a year, its [significant overstatements](#) of the time users spent watching videos on the platform (by 60 to 80 percent, Facebook says; by 150 to 900 percent, the plaintiffs say). [According to an exhaustive list at MarketingLand](#), over the past two years Facebook has admitted to misreporting the reach of posts on Facebook Pages (in two different ways), the rate at which viewers complete ad videos, the average time spent reading its "Instant Articles," the amount of referral traffic from Facebook to external websites, the number of views that videos received via Facebook's mobile site, and the number of video views in Instant Articles.

Can we still trust the metrics? After the Inversion, what's the point? Even when we put our faith in their accuracy, there's something not quite real about them: My favorite statistic this year was Facebook's claim that 75 million people watched at least a minute of Facebook Watch videos every day — [though, as Facebook admitted, the 60 seconds in that one minute didn't need to be watched consecutively](#). Real videos, real people, fake minutes.

The people are fake.

And maybe we shouldn't even assume that the people are real. Over at YouTube, the business of buying and selling video views is "flourishing," as [the Times reminded readers with a lengthy investigation in August](#). The company says only "a tiny fraction" of its traffic is fake, but fake subscribers are enough of a problem that the site undertook a purge of "spam accounts" in mid-December. These days, the *Times* found, you can buy 5,000 YouTube views — 30 seconds of a video counts as a view — for as low as \$15; oftentimes, customers are led to believe that the views they purchase come from real people. More likely, they come from bots. On some platforms, video views and app downloads can be forged in lucrative industrial counterfeiting operations. If you want a picture of what the Inversion looks like, [find a video of a "click farm"](#): hundreds of individual smartphones, arranged in rows on shelves or racks in professional-looking offices, each watching the same video or downloading the same app.

I never tire of looking at videos of Chinese click farms. It's just so surreal to see hundreds of phones playing the same video for the purposes of fake engagement.

pic.twitter.com/bHAGLqRqVb

— Matthew Brennan (@mbrennanchina) [December 10, 2018](#)

This is obviously not real human traffic. But what would real human traffic look like? The Inversion gives rise to some odd philosophical quandaries: If [a Russian troll using a Brazilian man's photograph to masquerade as an American Trump](#)

[supporter](#) watches a video on Facebook, is that view “real”? Not only do we have bots masquerading as humans and humans masquerading as other humans, but also sometimes humans masquerading as bots, pretending to be “[artificial-intelligence personal assistants](#),” like Facebook’s “M,” in order to help tech companies appear to possess cutting-edge AI. We even have [whatever CGI Instagram influencer Lil Miquela is](#): a fake human with a real body, a fake face, and real influence. Even humans who aren’t masquerading can contort themselves through layers of diminishing reality: [The Atlantic reports that non-CGI human influencers are posting fake sponsored content](#) — that is, content meant to look like content that is meant to look authentic, for free — to attract attention from brand reps, who, they hope, will pay them real money.

The businesses are fake.

The money is usually real. Not always — ask someone who enthusiastically got into cryptocurrency this time last year — but often enough to be an engine of the Inversion. If the money is real, why does anything else need to be? Earlier this year, the writer and artist Jenny Odell began to look into an Amazon reseller that had bought goods from other Amazon resellers and resold them, again on Amazon, at higher prices. [Odell discovered an elaborate network of fake price-gouging and copyright-stealing businesses connected to the cultlike Evangelical church](#) whose followers resurrected *Newsweek* in 2013 as a zombie search-engine-optimized spam farm. She visited a strange bookstore operated by the resellers in San Francisco and found a stunted concrete reproduction of the dazzlingly phony storefronts she'd encountered on Amazon, arranged haphazardly with best-selling books, plastic tchotchkes, and beauty products apparently bought from wholesalers. "At some point I began to feel like I was in a dream," she wrote. "Or that I was half-awake, unable to distinguish the virtual from the real, the local from the global, a product from a Photoshop image, the sincere from the insincere."

The content is fake.

The only site that gives me that dizzying sensation of unreality as often as Amazon does is YouTube, which plays host to weeks' worth of inverted, inhuman content. TV episodes that have been mirror-flipped to avoid copyright takedowns air next to huckster vloggers flogging merch who air next to [anonymously produced videos that are ostensibly for children](#). An animated video of Spider-Man and Elsa from *Frozen* riding tractors is not, you know, *not* real: Some poor soul animated it and gave voice to its actors, and I have no doubt that some number (dozens? Hundreds? Millions? Sure, why not?) of kids have sat and watched it and found some mystifying, occult enjoyment in it. But it's certainly not "official," and it's hard, watching it onscreen as an adult, to understand where it came from and what it means that the view count beneath it is continually ticking up.

These, at least, are mostly bootleg videos of popular fictional characters, i.e., counterfeit unreality. Counterfeit *reality* is still more difficult to find—for now. In January 2018, an anonymous Redditor created a relatively easy-to-use desktop-app implementation of "deepfakes," the now-infamous technology that uses artificial-intelligence image processing to replace one face in a video with another — putting, say, a politician's over a porn star's. [A recent academic paper from researchers at the graphics-card company Nvidia](#) demonstrates a similar technique used to create images of computer-generated "human" faces that look shockingly like photographs of real people. (Next time Russians want to puppeteer a group of invented Americans on Facebook, they won't even need to steal photos of real people.) Contrary to what you might expect, a world suffused with

deepfakes and other artificially generated photographic images won't be one in which "fake" images are routinely believed to be real, but one in which "real" images are routinely believed to be fake — simply because, in the wake of the Inversion, who'll be able to tell the difference?

Our politics are fake.

Such a loss of any anchoring “reality” only makes us pine for it more. Our politics have been inverted along with everything else, suffused with a Gnostic sense that we’re being scammed and defrauded and lied to but that a “real truth” still lurks somewhere. Adolescents are deeply engaged by YouTube videos that promise to show the hard reality beneath the “scams” of feminism and diversity — a process they call “red-pilling” after the scene in *The Matrix* when the computer simulation falls away and reality appears. Political arguments now involve trading accusations of “virtue signaling” — the idea that liberals are faking their politics for social reward — against charges of being Russian bots. The only thing anyone can agree on is that everyone online is lying and fake.

We ourselves are fake.

Which, well. Everywhere I went online this year, I was asked to prove I'm a human. Can you retype this distorted word? Can you transcribe this house number? Can you select the images that contain a motorcycle? I found myself prostrate daily at the feet of robot bouncers, frantically showing off my highly developed pattern-matching skills — does a Vespa count as a motorcycle, even? — so I could get into nightclubs I'm not even sure I want to enter. Once inside, I was directed by dopamine-feedback loops to scroll well past any healthy point, manipulated by emotionally charged headlines and posts to click on things I didn't care about, and harried and hectored and sweet-talked into arguments and purchases and relationships so algorithmically determined it was hard to describe them as real.

Where does that leave us? I'm not sure the solution is to seek out some pre-Inversion authenticity — to red-pill ourselves back to "reality." What's gone from the internet, after all, isn't "truth," but trust: the sense that the people and things we encounter are what they represent themselves to be. Years of metrics-driven growth, lucrative manipulative systems, and unregulated platform marketplaces, have created an environment where it makes more sense to be fake online — to be disingenuous and cynical, to lie and cheat, to misrepresent and distort — than it does to be real. Fixing that would require cultural and political reform in Silicon Valley and around the world, but it's our only choice. Otherwise we'll all end up on the bot internet of fake people, fake clicks, fake sites, and fake computers, where the only real thing is the ads.