

What Is “Surveillance Capitalism?” And How Did It Hijack the Internet?

by [Augustine Fou](#)

Shoshana Zuboff's new book *The Age of Surveillance Capitalism* goes into gory details of how companies collect, use, buy and sell your data for profit, often without consent or even the consumer knowing it was happening, until disasters reveal some of the dark underbelly—like the Cambridge Analytica scandal. But, I'm a marketer, so I will focus on the subset of “surveillance marketing”—also known as “digital marketing”—where companies profit off of you, because they are set up to do so. Digital ad-tech companies were built to extract as much value as possible from the trust transaction that used to be the user going to a publisher's site that carries an advertiser's ad.

Surveillance Marketing Was Built on the Foundation of Three Myths

Digital marketing as we know it today can be traced all the way back to Chris Anderson's book *The Long Tail*, published in 2006. Before that, digital media was primarily purchased from large sites that had large human audiences. *The Long Tail* promulgated the idea that collectively a large number of small sites could rival the scale of a small number of large sites. This simple premise alone led digital marketing down a dark and dangerous path to

the hell we now know is surveillance marketing. But most marketers don't even know they are in this hell. They were looking for scale in digital—and they got it. They were looking for data in digital—and they got it. And, they were looking for more granular targeting in digital—and they got it. But how?

Herein lies the three myths: 1) the long tail, 2) behavioral targeting and 3) hypertargeting.

The Myth of the Long Tail

The long tail of sites were all those super tiny, niche sites that had niche content, that people would visit. The theory was that with adtech, marketers could reach any person at any time with the right ad on any site. The reality was that virtually unlimited numbers of long tail sites could be created to carry ads, using ad tech, to absorb as many ad dollars as possible, under the guise of the scale of the long tail. But the nature of real long-tail sites means their content is so niche that the number of humans who actually were interested in the content would be tiny. That detail was overlooked or deliberately ignored when marketers were looking for scale. Fake long-tail sites were created, and bots were used to generate traffic for such sites, so marketers had more ad impressions to buy. But the scalability of digital ad tech—that is, without the limits of the physical world—combined with the greed of its creators led to an explosion of supply (digital ad inventory) that far outstripped even the large increase in demand, as more dollars from traditional channels like TV shifted into digital. The economics of this has been observed as decreasing CPMs (supply outruns demand) as opposed to increasing CPMs (slowly increasing supply but quickly increasing demand).

The Myth of Behavioral Targeting

With the rise of so many sites, the idea of behavioral targeting was spawned. Previously, we knew only the rough demographics of users based on the TV shows they watched or the large sites that they visited—and media was purchased based on these general categories. The theory of behavioral targeting was that if adtech could track what sites users visited, what pages they looked at, and even what content was on the page, they could figure out who the users were and what they wanted to buy, even if they never logged in or provided any personally identifiable information. But, consider this. It is relatively straightforward to believe that a user who visited a bunch of sports sites and a bunch of men's magazine sites is likely to be male, and a user who visited a lingerie site and a feminine hygiene product site was likely to be female. They might even be able to deduce that a user who visited GrandCanyon.com, REI.com and SmokeyBear.com was an outdoor enthusiast. But, as the nature and content of sites becomes more diverse, the assumptions and algorithms used to approximate characteristics of the user become less and less accurate. In fact, a recent study of online identifiers sold by DMPs (data management providers) revealed that more than 80% of the records were designated as *both* male and female—that is, they couldn't even get gender right. And, oh by the way, all those other 500 variables on that user are super accurate. Not.

The Myth of Hypertargeting

Finally, with all this data collected about users, adtech promised marketers that they could hypertarget them—literally, the right ad to the right person, at the right time, on any site they happen

to be visiting at the time. Right? Wrong. Here's why. Imagine you start with an audience and you choose one targeting parameter—male vs female. Using round numbers, if you target only males, you cut that audience size in half. Then if you choose a second targeting parameter, like age range—assume you have five ranges and select one—you cut that by five. You're already down to 10% of the audience size you started with. Add just one more targeting parameter, and the subset of the audience that matches all three selected parameters may already be down to 1% of the original audience size. What if you go to five parameters, ten, 50, 300 or 500 parameters? What fraction of the original audience will match all of those? Right. It's tiny. But the more parameters used for targeting, the more the adtech companies charge. A Carnegie Mellon study has since quantified the impact: "buying targeted ads over untargeted ads can be 500% times as expensive [for the marketer, but] in absolute terms the increase in revenues was \$0.000008 per advertisement [for the publisher]." In other words, it dramatically increased the profits of the adtech middlemen, but it harmed both the marketer and the publisher.

Marketers continue to believe these myths and increase the ad budgets they spend on it; but as far back as 2010, groups like the EFF (Electronic Frontier Foundation) were already sounding the alarm about "tracking without consent". They released a demo called Panopticlick in 2010 to show consumers what and how much was being tracked, even if they deleted cookies—that is, the consumer has no choice or recourse.

Surveillance Marketing Profits off of Three "Yous"—the Consumer, the Publisher and the Marketer

In the case of surveillance marketing, the “you” goes beyond just the consumer. All three parties—the consumer, advertisers and publishers—lose in this equation. Here’s how.

The consumer pays with his or her privacy. You’ve likely heard the phrase “you are the product.” This refers to consumers who use free products and services like Gmail and Facebook. But even though the consumers don’t pay money for those services, they pay with their privacy. Their personal data and activities are tracked and used to support the surveillance marketing industrial complex. Consumers weren’t fully aware of the extent to which they were being surveilled and still don’t have any recourse. But the more data, the more profit for these adtech middlemen.

The marketer pays with wasted ad spend. But more data and more targeting and more long tail sites do not translate into more business outcomes for marketers—for the reasons stated above in the three myths. But beyond the fact that surveillance marketing doesn’t work, there are new harms introduced into the digital marketing ecosystem—that is, ad fraud—and facilitated by the tech. The belief in the long tail led to the easy creation and proliferation of fake sites to carry ads. It also led to fake users—bots that visited sites repeatedly to create ad inventory out of thin air. The bots also would visit selected sites to make themselves look like any behavioral segment marketers wanted to target. Without the limits of the physical world—finite TV ad slots, finite print pages for ads, finite number of billboards by the road—it was easy to “hyperscale” everything to drive revenue and profits for the adtech industrial complex. There aren’t enough humans on earth that will flock to all websites and all devices and all forms of media at all times to generate all that

supposed ad inventory. But venture capitalists insist on it, so they can exit at hyper multiples. For marketers, it would have been better if those budgets were not spent on fraud and on the myths that didn't actually drive incremental business.

The publisher pays with lost ad revenue, or death. What do you think happens to those good publishers who created real content and, therefore, had real human audiences? They have real writers, journalists and editors, and creating original content is hard and expensive. They can't compete against fake sites that plagiarize everything or just make up everything—fake news, fake content and so on. Not only are CPMs pressured downward, but ad revenue actually flows away from good publishers to other sources, when marketers and their media buyers chase low-cost inventory on ad exchanges. The adtech industrial complex profits from this, and it helps fraudsters profit from it as well. Specifically, operators of fake sites that use fake traffic from fake users also can easily plug into the spigot of ad dollars by using adtech—whatever they siphon, the platforms get a cut too. In the good old days of digital, a marketer paid a publisher to run ads on its site, so the human audiences would see the ads. With the rise of adtech and the three myths, digital middlemen inserted themselves into the supply chain between the publisher and the marketer. The middlemen were all maximizing revenues and profits by extracting as much value as possible from the supply chain—to the point that good publishers could not pay their journalists and many are struggling to survive.

That's where the original "contract" of the internet was derailed.

Why Did It Go "Sideways"?

Previously, when people visited a website, they had a “first-party” interaction. They went to the site because they wanted to, and they got free content, because they understood the publisher made money by showing ads on the page. They even might have trusted the publisher and, therefore, the content on the site, because the user chose to go to the site. But when surveillance tech was added to the site, in the form of dozens upon dozens of third-party trackers, that one-to-one exchange between the user and the publisher became compromised. The user, and sometimes the publisher, didn’t know what data was being collected and sent off to third parties. Also, once the data was sent, neither the user or the publisher could do anything about it. That data became the “oil” that adtech traded on and profited from. See: [The Economist - The world’s most valuable resource is no longer oil, but data](#). But the publishers probably didn’t mean to violate the privacy of the people who trusted them; they simply lost control because third parties came in and did what they wanted to. Consumers’ privacy was being violated without consent or recourse. Publishers revenues were siphoned away by middlemen. And marketers lost ad budgets to fraud and waste and lower effectiveness. Perhaps adtech deserves to be renamed the “Badtech Industrial Complex”.



Figure 1. First party interaction now subverted by unknown third-party trackers.



Figure 2. How much do they extract?



Figure 3. "Badtech harms all parties and profits only "Badtech"

What Can Be Done?

End surveillance marketing. Do so with a three-step process: 1) protect consumers (specifically, privacy), 2) protect publishers (reduce adtech) and 3) protect advertisers/marketers (from fraud and ad waste). The process to end "surveillance marketing" is simple, but it may not be easy, because too many incumbent forces are at work to preserve it, so the Badtech Industrial Complex can continue to rape and pillage from all three parties in digital marketing: consumers, publishers and advertisers.

Ending surveillance marketing is necessary for the future of not only digital marketing but also the future of humankind. As we evolve headlong into a fully connected world, surveillance is the default, until we change it. Will the value of human dignity and privacy be sold off for the digital pennies that profit only the Badtech Industrial Complex to the detriment of all three original parties to an internet "transaction"? Or will someone step up and restore the trust transaction that was the user going to a publisher's site that carries an advertiser's ad?