



How Tech Utopia Fostered Tyranny

*Authoritarians' love for digital technology is no fluke—
it's a product of Silicon Valley's "smart" paternalism*

Jon Askonas

The rumors spread like wildfire: Muslims were secretly lacing a Sri Lankan village's food with sterilization drugs. Soon, a video circulated that appeared to show a Muslim shopkeeper admitting to drugging his customers—he had misunderstood the question that was angrily put to him. Then all hell broke loose. Over a several-day span, dozens of mosques and Muslim-owned shops and homes were burned down across multiple towns. In one home, a young journalist was trapped, and perished.

Mob violence is an old phenomenon, but the tools encouraging it, in this case, were not. As the *New York Times* reported in April, the rumors were spread via Facebook, whose newsfeed algorithm prioritized high-engagement content, especially videos. "Designed to maximize user time on site," as the *Times* article describes, the newsfeed algorithm "promotes whatever wins the most attention. Posts that tap into negative, primal emotions like anger or fear, studies have found, produce the highest engagement, and so proliferate." On Facebook in Sri Lanka, posts with incendiary rumors had among the highest engagement rates, and so were among the most highly promoted content on the platform. Similar cases of mob violence have taken place in India, Myanmar, Mexico, and elsewhere, with misinformation spread mainly through Facebook and the messaging tool WhatsApp.

Jon Askonas is assistant professor of politics and a Fellow at the Center for the Study of Statesmanship at the Catholic University of America.

WINTER 2019 ~ 3

This is in spite of Facebook’s decision in January 2018 to tweak its algorithm, apparently to prevent the kind of manipulation we saw in the 2016 U.S. election, when posts and election ads originating from Russia reportedly showed up in newsfeeds of up to 126 million American Facebook users. The company explained that the changes to its algorithm will mean that newsfeeds will be “showing more posts from friends and family and updates that spark conversation,” and “less public content, including videos and other posts from publishers or businesses.” But these changes, which Facebook had tested out in countries like Sri Lanka in the previous year, may actually have exacerbated the problem—which is that incendiary content, when posted by friends and family, is guaranteed to “spark conversation” and therefore to be prioritized in newsfeeds. This is because “misinformation is almost always more interesting than the truth,” as Mathew Ingram provocatively put it in the *Columbia Journalism Review*.

How did we get here, from Facebook’s mission to “give people the power to build community and bring the world closer together”? Riot-inducing “fake news” and election meddling are obviously far from what its founders intended for the platform. Likewise, Google’s founders surely did not build their search engine with the intention of its being censored in China to suppress free speech, and yet, after years of refusing this demand from Chinese leadership, Google has recently relented rather than pull their search engine from China entirely. And YouTube’s creators surely did not intend their feature that promotes “trending” content to help clickbait conspiracy-theory videos go viral.

These outcomes—not merely unanticipated by the companies’ founders but outright opposed to their intentions—are not limited to social media. So far, Big Tech companies have presented issues of incitement, algorithmic radicalization, and “fake news” as merely bumps on the road of progress, glitches and bugs to be patched over. In fact, the problem goes deeper, to fundamental questions of human nature. Tools based on the premise that access to information will only enlighten us and social connectivity will only make us more humane have instead fanned conspiracy theories, information bubbles, and social fracture. A tech movement spurred by visions of libertarian empowerment and progressive uplift has instead fanned a global resurgence of populism and authoritarianism.

Despite the storm of criticism, Silicon Valley has still failed to recognize in these abuses a sharp rebuke of its sunny view of human nature. It remains naïvely blind to how its own aspirations for social engineering are on a spectrum with the tools’ “unintended” uses by authoritarian regimes and nefarious actors.

AI Persuasion

The digital utopian dream of our age looks something like the 2016 concept video created by a Google R&D lab for a never-released product called the Selfish Ledger. The video was obtained in May by *The Verge*, which described it as “an unsettling vision of Silicon Valley social engineering.” Borrowing from Richard Dawkins’s notion of the “selfish gene,” the Selfish Ledger would be a self-help product on steroids, combining Google’s cornucopia of personal data with artificial-intelligence tools whose sole aim was to help you meet your goals.

Want to lose weight? Google Maps might prioritize smoothie shops or salad places when you search for “fast food.” Want to reduce your carbon footprint? Google might help you find vacation options closer to home or prioritize locally grown foods in the groceries that Google Express delivers to your doorstep. When the program needs more information than Google’s data banks can provide, it might suggest you buy a sensor, such as an Internet-connected scale or Google’s new AI-powered wearable camera. Or, if the needed product is not on the market, it might even suggest a design and 3D-print it.

The program is “selfish” in that it stubbornly pursues the self-identified goal the user gives it. But, the video explains, further down the road “suggestions may be converted not by the user but by the ledger itself.” And beyond individual self-help, by surveilling users over space and time Google would develop a “species-level understanding of complex issues such as depression, health, and poverty.”

The idea, according to a lab spokesperson, was meant only as a “thought-experiment...to explore uncomfortable ideas and concepts in order to provoke discussion and debate.” But the slope from Google’s original product—the seemingly value-neutral search engine—to the social engine of the Selfish Ledger is slipperier than one might think. The video’s vision of a smart Big Brother follows quite naturally from the company’s founding mission “to organize the world’s information, making it universally accessible and useful.” As Adam White recently wrote in these pages (“Google.gov,” Spring 2018), “Google has always understood its ultimate project not as one of rote descriptive recall but of *informativeness* in the fullest sense.”

After plucking the low-hanging fruit of web search, Google’s engineers began creating predictive search technologies like “autocomplete” and search results tailored to individual users based on their search histories. But what we are searching for—what we desire—is often shaped by

what we are exposed to and what we believe others desire. And so predicting what is useful, however value-neutral this may sound, can shade into *deciding* what is useful, both to individual users and to groups, and thereby shaping what kinds of people we become, for both better and worse.

The moral nature of usefulness becomes even clearer when we consider that our own desires are often in conflict. Someone may say he wants to have a decent sleep schedule, and yet his desire to watch another YouTube video about “deep state” conspiracy theories may get the better of him. Which of these two conflicting desires is the truer one? What is useful in this case, and what is good for him? Is he searching for conspiracy theories to find the facts of the matter, or to get the informational equivalent of a hit of cocaine? Which is more useful? What we wish for ourselves is often not what we do; the problem, it seemed to Walker Percy, is that modern man above all wants to know *who he is and should be*.

YouTube’s recommendation feature has helped to radicalize users through feedback loops—not only, again, by helping clickbait conspiracy videos go viral, but also by enticing users to view more videos like the ones they’ve already looked at, thus encouraging the user merely intrigued by extremist ideas to become a true diehard. Yet this result is not a curious fluke of the preference-maximizing vision, but its inevitable fruition. As long as our desires are unsettled and malleable—as long as we are human—the engineering choices of Google and the rest must be as much acts of persuasion as of prediction.

California Streamin’

The digital mindset of precisely measuring, analyzing, and ever more efficiently fulfilling our individual desires is of course not unique to Google. It pervades all of the Big Tech companies whose products give them access to massive amounts of user data, including also Facebook, Microsoft, Amazon, and to some extent Apple. Each company was founded on a variation of the premise that providing more people with more information and better tools, and helping them connect with each other, would help them lead better, freer, richer lives.

This vision is best understood as a descendant of the California counterculture, another way of extending decentralized, bottom-up power to the people. The story is told in Fred Turner’s 2006 book *From Counterculture to Cyberculture: Stewart Brand, the Whole Earth Network, and the Rise of Digital Utopianism*. Turner writes that Stewart Brand, erstwhile editor of the counterculture magazine *Whole Earth Catalog*, “suggested

that computers might become a new LSD, a new small technology that could be used to open minds and reform society.” Indeed, Steve Jobs came up with the name “Apple Computing” from living in an acid-infused commune at an Oregon apple orchard.

Not coincidentally, the tech giants are now investing heavily in using artificial intelligence to provide customized user experiences—not the information that is most useful to people in general, but to individual users.* The AI assistant is the culmination of utopian aspiration and shareholder value, a kind of techno-savvy guardian angel that perfectly and mysteriously knows how to meet your requests and sort your infinitely scrollable feed of search results, products, and friend updates, just for you. In the process, these companies run headfirst into the impossibility of separating the supposedly value-neutral criterion of usefulness from the moral aims of personal and social transformation.

For at the foundation of the digital revolution there was a hidden tension. First through personal computing and then through the Internet, the revolutionaries offered, as Brand’s *Whole Earth Catalog* put it, “access to tools.” A precious few users today grasp and take advantage of the full promise of networked computers to build ever more useful applications and tools. Instead, the vast majority spend their time and resources on only a few functions on a few platforms, consuming entertainment, searching for information, connecting with friends, and buying products or services.

And while in theory there are more “choices” and “flexibility” available than ever, in practice these are winner-take-all platforms, with the default choices and settings dominating user behavior. Google can return tens of millions of results for a search, but most users won’t leave the first page. Essentially random suggestions to users can become self-fulfilling prophecies, as *Wired* reported of the obscure 1988 climbing memoir *Touching the Void*, which by 2004 had become a hit due to Amazon’s recommendation algorithm.

Moreover, because algorithms are subject to strategic manipulation and because they are attempting to provide results unique to you, the choices shaping these powerful defaults are necessarily hidden away by platforms demanding you simply trust them. Ever since its founding,

* One must qualify that much of what is today called “artificial intelligence” is little more than traditional regression analysis, the basic technique taught in introductory statistics courses, but on an unprecedented scale and presence in daily life. None of this technology approaches the conscious, adaptive, reflective capacities often associated with the term, the kind we would find in *2001: A Space Odyssey*’s HAL 9000 or *Star Trek*’s Mr. Data. The labeling of these techniques as “artificial intelligence” arises in part from the ideological aspirations of Silicon Valley and in part from its overhyped marketing, and so ought to be resisted. But for the sake of critique we will adopt it here.

Google has had to keep its search algorithm's specific preferences secret and constantly re-adjust them to foil enterprising marketers trying to boost their profits at the expense of what users actually want. Every other Big Tech company has followed suit. As results have become more personalized, it becomes increasingly difficult to specify why, exactly, your newsfeed might differ from a friend's; the complex math behind it creates a black box that is "optimized" for some indiscernible set of metrics. Tech companies demand you simply trust the choices they make about how they manipulate results.

Much of the politics of Silicon Valley is explained by this Promethean exchange: gifts of enlightenment and ease in exchange for some measure of awe, gratitude, and deference to the technocratic elite that manufactures them. Algorithmic utopianism is at once optimistic about human motives and desires and paternalistic about humans' cognitive ability to achieve their stated preferences in a maximally rational way. Humans, in other words, are mostly good and well-intentioned but dumb and ignorant. We rely on poor intuitions and bad heuristics, but we can overcome them through tech-supplied information and cognitive adjustment. Silicon Valley wants to debug humanity, one default choice at a time.

We can see the shift from "access to tools" to algorithmic utopianism in the unheralded, inexorable replacement of the "page" by the "feed." The web in its earliest days was "surfing." Users actively explored what was interesting to them, shifting from page to page via links and URLs. While certain homepages—such as AOL or Yahoo!—were important, they were curated by actual people and communities. Most devoted "webizens" spent comparatively little time on them, instead exploring the web based on memory, bookmarks, and interests. Each blog, news source, store, and forum had its own site. Where life on the Internet didn't follow traditional editorial curation, it was mostly a do-it-yourself affair: Creating tools that might show you what your friends were up to, gathering all the information you cared about in one place, or finding new sites were rudimentary and tedious activities.

The feed was the solution to the tedium of surfing the web, of always having to decide for yourself what to do next. Information would now come to *you*. Gradually, the number of sites involved in one's life online dwindled, and the "platform" emerged, characterized by an infinite display of relevant information—the feed. The first feeds used fairly simple algorithms, but the algorithms have grown vastly more complex and personalized over time. These satisfaction-fulfillment machines are designed to bring you the most "relevant" content, where relevancy is ultimately

based on an elaborate and opaque model of who you are and what you want. But the opacity of these models, indeed the very personalization of them, means that a strong element of faith is required. By consuming what the algorithm says I want, I trust the algorithm to make me ever more who it thinks I already am.

In this process, users have gone from active surfers to sheep feeding at the algorithmic trough. Over time, platforms have come up with ever more sophisticated means of inducing behavior, both online and in real life, using AI-fueled notifications, messages, and default choices to nudge you in the right direction, ostensibly toward your own maximum satisfaction. Yet now, in order to rein in the bad behaviors the feeds themselves have encouraged—fake news, trolling, and so on—these algorithms have increasingly become the sites of stealthy intervention, using tweaks like “shadowbanning,” “down-ranking,” and simple erasure or blocking of users to help determine what information people do and don’t access, and thereby to subtly shape their minds.

Facebook in the Wild

Big Tech companies have thus married a fundamentally expansionary approach to information-gathering to a woeful naïveté about the likely uses of that technology. Motivated by left-liberal utopian beliefs about human progress, they are building technologies that are easily, naturally put to authoritarian and dystopian ends. While the Mark Zuckerbergs and Sergey Brins of the world claim to be shocked by the “abuse” of their platforms, the softly progressive ambitions of Silicon Valley and the more expansive visions of would-be dictators exist on the same spectrum of invasiveness and manipulation. There’s a sense in which the authoritarians have a better idea of what this technology is for.

Wasn’t it rosy to assume that the main uses of the most comprehensive, pervasive, automated surveillance and behavioral-modification technology in human history would be reducing people’s carbon footprints and helping them make better-informed choices in city council races? It ought to have been obvious that the new panopticon would be as liable to cut with the grain as against it, to become in the wrong hands a tool not for ameliorating but exploiting man’s natural capacity for error. Of the two sides, cheer for Dr. Jekyll, but bet on Mr. Hyde.

In recent years, two related problems have been shattering Silicon Valley’s dreams of progress. The first problem is that people have stubbornly refused to be debugged and empowered. Google hoped to provide

users with more “useful” information, but if you already know what you want to believe, Google exaggerates confirmation bias by feeding you more of what you want to hear. Facebook wanted to help people connect with their friends, share experiences, and learn from each other, but it turns out that people often pick the friends they want to engage with based on whether they care about the same things, leading the newsfeed algorithm to produce a custom-built echo chamber. Amazon stocks a wider selection of books than any store in history, but suggests them to you based on your search history and previous purchases, eliminating the cultivated, mind-broadening randomness of the bookstore browse.

In a sense, people often use these technologies backwards from how they were intended. In each case, what at first blush seems like a great tool for building what sociologists call “bridging capital”—connections to our neighbors or people in different interest groups—has in fact done far more to build “bonding capital”—tighter interconnections with people who are already like us in important ways.

This gap between what these systems are for and how they are actually used is amplified by globalization. Big Tech, to use a term from psychological research, is “WEIRD”—Western, educated, industrialized, rich, and democratic. These products were initially built by and for college-educated, Western, urban users. Facebook, for example, helped earn its early cachet by being exclusively for Harvard students (before it was expanded to Stanford, Columbia, and Yale). This means that the design choices product engineers make, and the behaviors those choices are designed to elicit, are often intended for a much more limited set of users than the technology will encounter “in the wild.”

A London economist, an underemployed Brazilian, and a Pakistani shepherd might each respond to the same algorithmic design choices with vastly different behaviors—in both the digital and the real world. Each of these big systems is designed, in its own way, to maximize user engagement, but what content users engage with, and how, depends in large part on culture, class, and psychology.

For a WEIRD user working in journalism or politics, “user engagement” might mean an addiction to Twitter. For a teenage girl on Instagram, it might lead to anorexia and depression. Among Sri Lankan villagers, it was a recipe for “fake news,” overheated rhetoric, and riotous violence. As the *New York Times* article on the story explained, “Online outrage mobs will be familiar to any social media user. But in places with histories of vigilantism, they can work themselves up to real-world attacks.”

These technologies were based on a model in which users' desires were crafted outside the system, and the purpose of the algorithms was to measure and meet those desires with ever greater efficiency. The designers did not imagine the algorithms themselves shaping users by feeding their basest impulses, turning the high of a notification ping into whatever behaviors result in more pings—snarkier tweets, sexier pictures, or more feverish posts. The engineering choices that have made these technologies so compelling and addictive have also made it completely implausible that they would fulfill their founders' noble ambitions. Like Dr. Frankenstein, Big Tech's creators in no way control their creations.

Surveillance State, Made in U.S.A.

Thus we arrive at the second problem besetting Big Tech: Malicious actors, authoritarian regimes chief among them, are sophisticated adopters and promoters of the information revolution. How long ago were the halcyon days of the Arab Spring, when commentators could argue that Facebook and Twitter presented an existential threat to dictatorships everywhere. In reality, authoritarian regimes the world over quickly learned to love technologies that enticed their subjects into carrying around listening devices and putting their innermost thoughts online.

Big Brother can read tweets too, which is why China's massive surveillance system includes monitoring social media. Slowing down Internet traffic, as Iran has apparently done, turns out to be an even more effective source of censorship than outright blocking of websites—accessing information becomes a matter of great frustration instead of forbidden allure. Before Russian troll farms were aimed at American Facebook users, they were found to be useful at home for stirring up anti-American sentiments and defending Russia's aggressions in Ukraine.

By pulling so much of social life into cyberspace, the information revolution has made dissent more visible, manageable, and manipulable than ever before. Hidden public anger, the ultimate *bête noire* of many a dictator, becomes more legible to the regime. Activating one's own supporters, and manipulating the national conversation, become easier as well. Indeed, the information revolution has been a boon to the police state. It used to be incredibly manpower-intensive to monitor videos, accurately take and categorize images, analyze opposition magazines, track the locations of dissidents, and appropriately penalize enemies of the regime. But now, tools that were perfected for tagging your friends in beach photos, categorizing new stories, and ranking products by user reviews are the

technological building blocks of efficient surveillance systems. Moreover, with big data and AI, regimes can now engage in especially “smart” forms of what is sometimes called “smart repression”—exerting just the right amount of force and nudging, at the lowest possible cost, to pull subjects into line. The computational counterculture’s promise of “access to tools” and “people power” has, paradoxically, contributed to mass surveillance and oppression.

What’s shocking isn’t that technological development is a two-edged sword. It’s that the power of these technologies is paired with a stunning apathy among their creators about who might use them and how. Google employees have recently declared that helping the Pentagon with a military AI program is a bridge too far, convincing the company to cancel a \$10 billion contract. But at the same time, Google, Apple, and Microsoft, committed to the ideals of open-source software and collaboration toward technological progress, have published machine-learning tools for anyone to use, including agents provocateur and revenge pornographers.

In 2017 researchers from the tech company Nvidia published an algorithm for realistically modifying video, for example to turn a winter scene into a summer scene. Within months, as *Motherboard* reported, an anonymous Internet hobbyist had developed a similar technology to create and release software for swapping faces in videos with high fidelity. While the intent was (inevitably) pornographic, the political implications of the technology were immediately recognized, as in a BuzzFeed video of a fake announcement by former President Obama. Recently, IBM announced the creation of a free database of over one million racially diverse facial images to help train facial recognition algorithms and reduce bias. One wonders whether the Uighur people arrested by the Chinese government with the help of facial recognition technology are grateful that they weren’t discriminated against.

Silicon Valley’s tech founders envisioned a world where information technology directly contributed to an increasingly democratic society, characterized by decentralization, a do-it-yourself attitude, and an independence of thought associated with both their brand of Sixties counterculture and a deeper American tradition. They and their successors, based on optimistic assumptions about human nature, built machines to maximize those naturally good human desires. But, to use a line from Bruno Latour, “technology is society made durable.” That is, to extend Latour’s point, technology stabilizes in concrete form what societies already find desirable.

The counterculture’s humanism has long been overthrown by dreams of maximizing satisfaction, metrics, profits, “knowledge,” and connection,

a task now to be given over to the machines. The emerging soft authoritarianism in Silicon Valley's designs to stoke our desires will go hand in hand with a hard authoritarianism that pushes these technologies toward their true ends.