

[Facebook Tried to Make Its Platform a Healthier Place. It Got Angrier Instead.](#)

By Keach Hagey and Jeff Horwitz

In the fall of 2018, Jonah Peretti, chief executive of online publisher BuzzFeed, emailed a top official at Facebook Inc. The most divisive content that publishers produced was going viral on the platform, he said, creating an incentive to produce more of it.

He pointed to the success of a BuzzFeed post titled “21 Things That Almost All White People are Guilty of Saying,” which received 13,000 shares and 16,000 comments on Facebook, many from people criticizing BuzzFeed for writing it, and arguing with each other about race. Other content the company produced, from news videos to articles on self-care and animals, had trouble breaking through, he said.

Mr. Peretti blamed a major overhaul Facebook had given to its News Feed algorithm earlier that year to boost “meaningful social interactions,” or MSI, between friends and family, according to internal Facebook documents reviewed by The Wall Street Journal that quote the email.

BuzzFeed built its business on making content that would go viral on Facebook and other social media, so it had a vested interest in any algorithm changes that hurt its distribution. Still, Mr. Peretti’s email touched a nerve.

Facebook’s chief executive, Mark Zuckerberg, said the aim of the algorithm change was to strengthen bonds between users and to improve their well-being. Facebook would encourage people to interact more with friends and family and spend less time passively consuming professionally produced content, which research suggested was harmful to their mental health.

Within the company, though, staffers warned the change was having the opposite effect, the documents show. It was making Facebook’s platform an angrier place.

Company researchers discovered that publishers and political parties were reorienting their posts toward outrage and sensationalism. That tactic produced high levels of comments and reactions that translated into success on Facebook.

“Our approach has had unhealthy side effects on important slices of public content, such as politics and news,” wrote a team of data scientists, flagging Mr. Peretti’s complaints, in a memo reviewed by the Journal. “This is an increasing liability,” one of them wrote in a later memo.

They concluded that the new algorithm’s heavy weighting of reshared material in its News Feed made the angry voices louder. “Misinformation, toxicity, and violent content are inordinately prevalent among reshares,” researchers noted in internal memos.

Some political parties in Europe told Facebook the algorithm had made them shift their policy positions so they resonated more on the platform, according to the documents.

“Many parties, including those that have shifted to the negative, worry about the long term effects on democracy,” read one internal Facebook report, which didn’t name specific parties.

Facebook employees also discussed the company’s other, less publicized motive for making the change: Users had begun to interact less with the platform, a worrisome trend, the documents show.

The email and memos are part of an extensive array of internal company communications reviewed by the Journal. They offer an unparalleled look at how much Facebook knows about the flaws in its platform and how it often lacks the will or the ability to address them. This is the third in a series of articles based on that information.

In an interview, Lars Backstrom, a Facebook vice president of engineering, said that any algorithm risks promoting content that is objectionable or harmful to some users.

“Like any optimization, there’s going to be some ways that it gets exploited or taken advantage of,” he said. “That’s why we have an integrity team that is trying to track those down and figure out how to mitigate them as efficiently as possible.”

Data scientists on that integrity team—whose job is to improve the quality and trustworthiness of content on the platform—worked on a number of potential changes to curb the tendency of the overhauled algorithm to reward outrage and lies. Mr. Zuckerberg resisted some of the proposed fixes, the documents show, because he was worried they might hurt the company’s other objective—making users engage more with Facebook.

Anna Stepanov, who led a team addressing those issues, presented Mr. Zuckerberg with several proposed changes meant to address the proliferation of false and divisive content on the platform, according to an April 2020 internal memo she wrote about the briefing. One such change would have taken away a boost the algorithm gave to content most likely to be reshared by long chains of users.

“Mark doesn’t think we could go broad” with the change, she wrote to colleagues after the meeting. Mr. Zuckerberg said he was open to testing the approach, she said, but “We wouldn’t launch if there was a material tradeoff with MSI impact.”

Last month, nearly a year and a half after Ms. Stepanov said Mr. Zuckerberg nixed the idea of broadly incorporating a similar fix, Facebook announced it was “gradually expanding some tests to put less emphasis on signals such as how likely someone is to comment or share political content.” The move is part of a broader push, spurred by user surveys, to reduce the amount of political content on Facebook after the company came under criticism for the way election protesters used the platform to question the results and organize protests that led to the Jan. 6 riot at the Capitol in Washington.

Mr. Backstrom, who oversees content ranking in News Feed, said Facebook made the recent change because it felt that the downsides of relying on engagement-based metrics for sensitive content categories such as politics outweighed the benefits.

The 2018 algorithm change affected Facebook's central feature, the News Feed, a constantly updated, personally customized scroll of friends' family photos and links to news stories. It accounts for the majority of time Facebook's nearly three billion users spend on the platform. The company sells that user attention to advertisers, both on Facebook and its sister platform Instagram, accounting for nearly all of its \$86 billion in revenue last year.

A proprietary algorithm controls what appears in each user's News Feed. It takes into account who users are friends with, what kind of groups they have joined, what pages they have liked, which advertisers have paid to target them and what types of stories are popular or driving conversation.

Significant changes to the algorithm can have major implications for the company, advertisers and publishers. Facebook has made many algorithm tweaks over the years. The shift to emphasize MSI was one of the biggest.

"Is a ranking change the source of the world's divisions? No," said Facebook spokesman Andy Stone in a written statement. "Research shows certain partisan divisions in our society have been growing for many decades, long before platforms like Facebook even existed."

News Corp, owner of The Wall Street Journal, has a commercial agreement to supply news through Facebook.

'The right thing'

In January 2018, Facebook was coming off a trying year. It was on the defensive in Washington about what U.S. intelligence officials said was Russia's use of the platform to meddle in the 2016 U.S. presidential election.

Mr. Zuckerberg announced he was changing Facebook product managers' goal from helping people find relevant content to helping them interact more with friends and family.

He said the shift was driven by research showing that passive media consumption on Facebook—notably video, which had been exploding on the platform—wasn't as good for well-being as interacting with other people.

He framed the change as a sacrifice. "Now, I want to be clear: by making these changes, I expect the time people spend on Facebook and some measures of engagement will go down," he wrote on Facebook. "But I also expect the time you do spend on Facebook will be more valuable. And if we do the right thing, I believe that will be good for our community and our business over the long term too."

Facebook training videos and internal memos show another reason for the change—the company's growing concern about a decline in user engagement, which typically refers to actions like commenting on or sharing posts. Engagement is viewed inside the company as an important sign for the health of the business.

Comments, likes and reshares declined through 2017, while “original broadcast” posts—the paragraph and photo a person might post when a dog dies—continued a yearslong decline that no intervention seemed able to stop, according to the internal memos. The fear was that eventually users might stop using Facebook altogether.

One data scientist said in a 2020 memo that Facebook teams studied the issue and “never really figured out why metrics declined.” The team members ultimately concluded that the prevalence of video and other professionally produced content, rather than organic posts from individuals, was likely part of the problem.

The goal of the algorithm change was to reverse the decline in comments, and other forms of engagement, and to encourage more original posting. It would reward posts that garnered more comments and emotion emojis, which were viewed as more meaningful than likes, the documents show.

In an internal training video, one Facebook employee said that in addition to the company’s “ethical duty” not to turn users into zombies with too much video, it had business reasons for intervening.

“People will probably leave the app if it’s bad for them,” the employee said.

Ranking reshares

Facebook’s solution was to create a formula that measured how much “meaningful” interaction a post sparked, then organize the News Feed to encourage as much of that as possible. Under an internal point system used to measure its success, a “like” was worth one point; a reaction, reshare without text or reply to an invite was worth five points; and a significant comment, message, reshare or RSVP, 30 points. Additional multipliers were added depending on whether the interaction was between members of a group, friends or strangers.

From a business perspective, it worked. As Facebook predicted, time spent on the platform declined, but the effort to maximize interactions between users slowed the free fall in comments, and mostly improved the all-important metric of “daily active people” using Facebook, according to tests run in August 2018, internal memos show.

The change had some positive effects. Content shared by close connections was more trustworthy, and users found it more meaningful than material from more distant acquaintances, according to one memo. But those benefits were outweighed by the aspect of the algorithm change that favored user actions like reshares and comments.

In an early sign of trouble, during the summer of 2018, Facebook data scientists repeatedly surveyed users and found that many felt the quality of their feeds had decreased, the documents show.

As Facebook had warned was likely, the change hurt many online publishers. In the first half of 2018, BuzzFeed suffered a 13% decline in traffic compared with the prior six months, Breitbart lost 46% and ABC News lost 12%, according to online data firm Comscore.

Topix, a news site known for local community forums, lost 53% of its traffic over that period. “Facebook was our major traffic driver at this time,” said Chris Tolles, Topix’s former CEO. “That was certainly a problem for us.” Topix was purchased by Publishers Clearing House in 2019.

In interviews, more than a dozen publishing executives described Facebook’s shift as the final straw after several years of constantly changing policies that convinced them they couldn’t rely on Facebook for traffic.

BuzzFeed rose to prominence on internet confections like lists and quizzes, but over the years built a robust news operation and expanded aggressively into video. Facebook has historically been a major source of traffic for BuzzFeed, though the publisher has diversified its distribution in recent years.

Buzzfeed’s Mr. Peretti, in his email, wrote that the new algorithm seemed to be disproportionately rewarding divisiveness, based on what the publisher saw in its own numbers and his observations about how other publishers’ posts performed.

“MSI ranking isn’t actually rewarding content that drives meaningful social interactions,” Mr. Peretti wrote in his email to the Facebook official, adding that his staff felt “pressure to make bad content or underperform.”

It wasn’t just material that exploited racial divisions, he wrote, but also “fad/junky science,” “extremely disturbing news” and gross images.

Political effect

In Poland, the changes made political debate on the platform nastier, Polish political parties told the company, according to the documents. The documents don’t specify which parties.

“One party’s social media management team estimates that they have shifted the proportion of their posts from 50/50 positive/negative to 80% negative, explicitly as a function of the change to the algorithm,” wrote two Facebook researchers in an April 2019 internal report.

Nina Jankowicz, who studies social media and democracy in Central and Eastern Europe as a fellow at the Woodrow Wilson Center in Washington, said she has heard complaints from many political parties in that region that the algorithm change made direct communication with their supporters through Facebook pages more difficult. They now have an incentive, she said, to create posts that rack up comments and shares—often by tapping into anger—to get exposure in users’ feeds.

The Facebook researchers, wrote in their report that in Spain, political parties run sophisticated operations to make Facebook posts travel as far and fast as possible.

“They have learnt that harsh attacks on their opponents net the highest engagement,” they wrote. “They claim that they ‘try not to,’ but ultimately ‘you use what works.’ ”

In the 15 months following fall 2017 clashes in Spain over Catalan separatism, the percentage of insults and threats on public Facebook pages related to social and political debate in Spain increased by 43%, according to research conducted by Constella Intelligence, a digital risk protection firm.

Facebook researchers wrote in their internal report that they heard similar complaints from parties in Taiwan and India.

Brad Parscale, who was the digital strategy leader for Donald Trump’s 2016 presidential campaign, and boasted that Facebook is where Mr. Trump won the election, said he began to notice changes to the algorithm as early as mid-2017, when the performance of political videos began to decline.

“ ‘Healthy’ is a cover word for ‘We need to reduce conservative and Trump video spread,’ ” he said in an interview for this article.

Republican National Committee Chair Ronna McDaniel, then-House Majority Leader Kevin McCarthy (R., Calif.) and Mr. Parscale met with Facebook executives in June 2018 to register their complaints. At that meeting, which included Facebook global public-policy chief Joel Kaplan, the Republicans were told the algorithm changes were meant to counter the spread of what the company considered misinformation, according to Mr. Parscale.

Proposed fixes

In April 2019, one Facebook data scientist proposed reducing the spread of “deep reshares,” which means the viewer is not a friend or follower of the original poster, according to an internal memo.

“While the FB platform offers people the opportunity to connect, share and engage, an unfortunate side effect is that harmful and misinformative content can go viral, often before we can catch it and mitigate its effects,” he wrote. “Political operatives and publishers tell us that they rely more on negativity and sensationalism for distribution due to recent algorithmic changes that favor reshares.”

Later, Facebook data scientists zeroed in on an aspect of the revamped algorithm called “downstream MSI,” which made a post more likely to appear in a user’s News Feed if the algorithm calculated people were likely to share or comment on it as it passed down the chain of reshares.

Early tests showed how reducing that aspect of the algorithm for civic and health information helped reduce the proliferation of false content. Facebook made the change for those categories in the spring of 2020.

When Ms. Stepanov presented Mr. Zuckerberg with the integrity team’s proposal to expand that change beyond civic and health content—and a few countries such as Ethiopia and Myanmar where changes

were already being made—Mr. Zuckerberg said he didn't want to pursue it if it reduced user engagement, according to the documents.

James Barnes, a former Facebook employee who left in 2019, said Facebook had hoped that giving priority to user engagement in the News Feed would bring people closer together. But the platform had grown so complex the company didn't understand how the change might backfire.

"There are no easy silver bullets," he said.