

Facebook Struggles to Balance Civility and Growth

Employees and executives are battling over how to reduce misinformation and hate speech without hurting the company's bottom line.



By **Kevin Roose, Mike Isaac and Sheera Frenkel**

Published Nov. 24, 2020 Updated Jan. 7, 2021

SAN FRANCISCO — In the tense days after the presidential election, a team of Facebook employees presented the chief executive, Mark Zuckerberg, with an alarming finding: Election-related misinformation was going viral on the site.

President Trump was already casting the election as rigged, and stories from right-wing media outlets with false and misleading claims about discarded ballots, miscounted votes and skewed tallies were among the most popular news stories on the platform.

In response, the employees proposed an emergency change to the site's news feed algorithm, which helps determine what more than two billion people see every day. It involved emphasizing the importance of what Facebook calls "news ecosystem quality" scores, or N.E.Q., a secret internal ranking it assigns to news publishers based on signals about the quality of their journalism.

Typically, N.E.Q. scores play a minor role in determining what appears on users' feeds. But several days after the election, Mr. Zuckerberg agreed to increase the weight that Facebook's algorithm gave to N.E.Q. scores to make sure authoritative news appeared more prominently, said three people with knowledge of the decision, who were not authorized to discuss internal deliberations.

The change was part of the "break glass" plans Facebook had spent months developing for the aftermath of a contested election. It resulted in a spike in visibility for big, mainstream publishers like CNN, The New York Times and NPR, while posts from highly engaged hyperpartisan pages, such as Breitbart and Occupy Democrats, became less visible, the employees said.

It was a vision of what a calmer, less divisive Facebook might look like. Some employees argued the change should become permanent, even if it was unclear how that might affect the amount of time people spent on Facebook. In an employee meeting the week after the

election, workers asked whether the “nicer news feed” could stay, said two people who attended.

Guy Rosen, a Facebook executive who oversees the integrity division that is in charge of cleaning up the platform, said on a call with reporters last week that the changes were always meant to be temporary. “There has never been a plan to make these permanent,” he said. John Hegeman, who oversees the news feed, said in an interview that while Facebook might roll back these experiments, it would study and learn from them.

The news feed debate illustrates a central tension that some inside Facebook are feeling acutely these days: that the company’s aspirations of improving the world are often at odds with its desire for dominance.

In the past several months, as Facebook has come under more scrutiny for its role in amplifying false and divisive information, its employees have clashed over the company’s future. On one side are idealists, including many rank-and-file workers and some executives, who want to do more to limit misinformation and polarizing content. On the other side are pragmatists who fear those measures could hurt Facebook’s growth, or provoke a political backlash that leads to painful regulation.

“There are tensions in virtually every product decision we make and we’ve developed a companywide framework called ‘Better Decisions’ to ensure we make our decisions accurately, and that our goals are directly connected to delivering the best possible experiences for people,” said Joe Osborne, a Facebook spokesman.

These battles have taken a toll on morale. In an employee survey this month, Facebook workers reported feeling less pride in the company compared to previous years. About half felt that Facebook was having a positive impact on the world, down from roughly three-quarters earlier this year, according to a copy of the survey, known as Pulse, which was reviewed by The New York Times. Employees’ “intent to stay” also dropped, as did confidence in leadership.

BuzzFeed News previously reported on the survey results.

Even as Election Day and its aftermath have passed with few incidents, some disillusioned employees have quit, saying they could no longer stomach working for a company whose products they considered harmful. Others have stayed, reasoning they can make more of a difference on the inside. Still others have made the moral calculation that even with its flaws, Facebook is, on balance, doing more good than harm.

“Facebook salaries are among the highest in tech right now, and when you’re walking home with a giant paycheck every two weeks, you have to tell yourself that it’s for a good cause,” said Gregor Hochmuth, a former engineer with Instagram, which Facebook owns, who left

in 2014. “Otherwise, your job is truly no different from other industries that wreck the planet and pay their employees exorbitantly to help them forget.”

With most employees working remotely during the pandemic, much of the soul-searching has taken place on Facebook’s internal Workplace network.

In May, during the heat of the Black Lives Matter protests, Mr. Zuckerberg angered many employees when he declined to remove a post by President Trump that said “when the looting starts, the shooting starts.” Lawmakers and civil rights groups said the post threatened violence against protesters and called for it to be taken down. But Mr. Zuckerberg said the post did not violate Facebook’s rules.

To signal their dissatisfaction, several employees formed a new Workplace group called “Take Action.” People in the group, which swelled to more than 1,500 members, pointedly changed their profile photos to an image of a raised “Black Lives Matter” fist.

The group became a home for internal dissent and dark humor about Facebook’s foibles. On several occasions, employees reacted to negative news stories about the company by posting a meme from a British comedy sketch in which two Nazis have a moral epiphany and ask themselves, “Are we the baddies?”

In June, employees staged a virtual walkout to protest Mr. Zuckerberg’s decisions regarding Mr. Trump’s posts.

In September, Facebook updated its employee policies to discourage workers from holding contentious political debates in open Workplace forums, saying they should confine the conversations to specifically designated spaces. It also required employees to use their real faces or the first initial of their names as their profile photo, a change interpreted by some workers as a crackdown.

Several employees said they were frustrated that to tackle thorny issues like misinformation, they often had to demonstrate that their proposed solutions wouldn’t anger powerful partisans or come at the expense of Facebook’s growth.

The trade-offs came into focus this month, when Facebook engineers and data scientists posted the results of a series of experiments called “P(Bad for the World).”

The company had surveyed users about whether certain posts they had seen were “good for the world” or “bad for the world.” They found that high-reach posts — posts seen by many users — were more likely to be considered “bad for the world,” a finding that some employees said alarmed them.

So the team trained a machine-learning algorithm to predict posts that users would consider “bad for the world” and demote them in news feeds. In early tests, the new algorithm successfully reduced the visibility of objectionable content. But it also lowered the number of times users opened Facebook, an internal metric known as “sessions” that executives monitor closely.

“The results were good except that it led to a decrease in sessions, which motivated us to try a different approach,” according to a summary of the results, which was posted to Facebook’s internal network and reviewed by The Times.

The team then ran a second experiment, tweaking the algorithm so that a larger set of “bad for the world” content would be demoted less strongly. While that left more objectionable posts in users’ feeds, it did not reduce their sessions or time spent.

That change was ultimately approved. But other features employees developed before the election never were.

One, called “correct the record,” would have retroactively notified users that they had interacted with false news and directed them to an independent fact-check. Facebook employees proposed expanding the product, which is currently used to notify people who have shared Covid-19 misinformation, to apply to other types of misinformation.

But that was vetoed by policy executives who feared it would disproportionately show notifications to people who shared false news from right-wing websites, according to two people familiar with the conversations.

Another product, an algorithm to classify and demote “hate bait” — posts that don’t strictly violate Facebook’s hate speech rules, but that provoke a flood of hateful comments — was limited to being used only on groups, rather than pages, after the policy team determined that it would primarily affect right-wing publishers if it were applied more broadly, said two people with knowledge of the conversations.

Mr. Rosen, the Facebook integrity executive, disputed those characterizations in an interview, which was held on the condition that he not be quoted directly.

He said that the “correct the record” tool wasn’t as effective as hoped, and that the company had decided to focus on other ways of curbing misinformation. He also said applying the “hate bait” detector to Facebook pages could unfairly punish publishers for hateful comments left by their followers, or make it possible for bad actors to hurt a page’s reach by spamming it with toxic comments. Neither project was shelved because of political concerns or because it reduced Facebook usage, he said.

“No News Feed product change is ever solely made because of its impact on time spent,” said Mr. Osborne, the Facebook spokesman. He added that the people talking to The Times had no decision-making authority.

Facebook's moves to clean up its platform will be made easier, in some ways, by the end of the Trump administration. For years, Mr. Trump and other leading conservatives accused the company of anti-conservative bias each time it took steps to limit misinformation.

But even with an incoming Biden administration, Facebook will need to balance employees' desire for social responsibility with its business goals.

“The question is, what have they learned from this election that should inform their policies in the future?” said Vanita Gupta, the chief executive of the civil rights group Leadership Conference on Civil and Human Rights. “My worry is that they’ll revert all of these changes despite the fact that the conditions that brought them forward are still with us.”

In a virtual employee meeting last week, executives described what they viewed as Facebook's election successes, said two people who attended. While the site was still filled with posts falsely claiming the election was rigged, Chris Cox, Facebook's chief product officer, said he was proud of how the company had applied labels to election-related misinformation, pointing users to authoritative information about the results, the people said.

Then the stream cut to a pre-produced video, a Thanksgiving morale-booster featuring a parade of employees talking about what they were grateful for this year.