

Who Controls Your Facebook Feed

A small team of engineers in Menlo Park. A panel of anonymous power users around the world. And, increasingly, you.

By Will Oremus



Photo illustration by Lisa Larson-Walker. Photo by Tang Ming Tung/Getty Images.

Every time you open Facebook, one of the world's most influential, controversial, and misunderstood algorithms springs into action. It scans and collects everything posted in the past week by each of your friends, everyone you follow, each group you belong to, and every Facebook page you've liked. For the average Facebook user, that's more than 1,500 posts. If you have several hundred friends, it could be as many as 10,000. Then, according to a closely guarded and constantly shifting formula, Facebook's news feed algorithm ranks them all, in what it believes to be the precise order of how likely you are to find each post worthwhile. Most users will only ever see the top few hundred.



WILLOREMUS

Will Oremus is *Slate's* senior technology writer. Email him at will.oremus@slate.com or follow him on [Twitter](#).

No one outside Facebook knows for sure how it does this, and no one inside the company will tell you. And yet the results of this automated ranking process shape the social lives and reading habits of more than 1 billion daily active users—one-fifth of the world's adult population. The algorithm's viral power has turned the media industry upside down, propelling startups like *BuzzFeed* and *Vox* to national prominence while 100-year-old newspapers wither and die. It fueled the stratospheric rise of billion-dollar companies like *Zynga* and *LivingSocial*—only to suck the helium from them a year or two later with a few adjustments to its code, leaving behind empty-pocketed investors and laid-off workers. Facebook's news feed algorithm can be tweaked to **make us happy or sad**; it can expose us to new and challenging ideas or insulate us in **ideological bubbles**.

Facebook through the Years

Facebook's news feed algorithm has shaped not only what we read and how we keep in touch, but how the media frame stories to catch our attention. From the start, savvy publishers have gamed the algorithm's quirks to concoct viral hits. In response, Facebook's engineers are constantly tweaking the code in ways that disadvantage some types of posts while boosting others. To see how the media's formulas for viral success have changed over the years, try to match each of these Facebook hits to the year it was published.

Interactive template by Chris Kirk. Report a bug or give feedback here.

Advertisement

And yet, for all its power, Facebook's news feed algorithm is surprisingly inelegant, maddeningly mercurial, and stubbornly opaque. It remains as likely as not to serve us posts we find trivial, irritating, misleading, or just plain boring. And Facebook knows it. Over the past several months, the social network has been running a test in which it shows some users the top post in their news feed alongside one other, lower-ranked post, asking them to pick the one they'd prefer to read. The result? The algorithm's rankings correspond to the user's preferences "sometimes," Facebook acknowledges, declining to get more specific. When they don't match up, the company says, that points to "an area for improvement."

"Sometimes" isn't the success rate you might expect for such a vaunted and feared bit of code. The news feed algorithm's outsized influence has given rise to a strand of criticism that treats it as if it possessed a mind of its own—as if it were some runic form of intelligence, loosed on the world to pursue ends beyond the ken of human understanding. At a time when Facebook and other Silicon Valley giants increasingly filter our choices and guide our decisions through **machine-learning software**, when tech titans like Elon Musk and scientific laureates like Stephen Hawking are warning of the **existential threat posed by A.I.**, the word itself—*algorithm*—has begun to take on an eerie affect. Algorithms, in the popular imagination, are mysterious, powerful entities that stand for all the ways technology and modernity both serve our every desire and threaten the values we hold dear.

The reality of Facebook's algorithm is somewhat less fantastical, but no less fascinating. I had a rare chance recently to spend time with Facebook's news feed team at their Menlo Park, California, headquarters and see what it actually looks like when they make one of those infamous, market-moving "tweaks" to the algorithm—why they do it, how they do it, and how they decide whether it worked. A glimpse into its inner workings sheds light not only on the mechanisms of Facebook's news feed, but on the limitations of machine learning, the pitfalls of data-driven decision making, and the moves Facebook is increasingly making to collect and address feedback from individual human users, including a growing panel of testers that are becoming Facebook's equivalent of the Nielsen family.

Facebook's algorithm, I learned, isn't flawed because of some glitch in the system. It's flawed because, unlike the perfectly realized, sentient algorithms of our **sci-fi fever dreams**, the intelligence behind Facebook's software is fundamentally human. Humans decide what data goes into it, what it can do with that data, and what they want to come out the other end. When the algorithm errs, humans are to blame. When it evolves, it's because a bunch of humans read a bunch of spreadsheets, held a bunch of meetings, ran a bunch of tests, and decided to make it better. And if it does keep getting better? That'll be because another group of humans keeps telling them about all the ways it's falling short: us.



When I arrive at Facebook’s sprawling, Frank Gehry–designed office in Menlo Park, I’m met by a lanky 37-year-old man whose boyish countenance shifts quickly between an earnest smile and an expression of intense focus. Tom Alison is director of engineering for the news feed; he’s in charge of the humans who are in charge of the algorithm.

Alison steers me through a maze of cubicles and open minikitchens toward a small conference room, where he promises to demystify the Facebook algorithm’s true nature. On the way there, I realize I need to use the bathroom and ask for directions. An involuntary grimace crosses his face before he apologizes, smiles, and says, “I’ll walk you there.” At first I think it’s because he doesn’t want me to get lost. But when I emerge from the bathroom, he’s still standing right outside, and it occurs to me that he’s not allowed to leave me unattended.



The Facebook office in Menlo Park, California.

Photo by Christophe Wu/Facebook

For the same reason—Facebook’s fierce protection of trade secrets—Alison cannot tell me much about the actual code that composes the news feed algorithm. He can, however, tell me what it does, and why—and why it’s always changing. He starts, as engineers often do, at the whiteboard.

“When you study computer science, one of the first algorithms you learn is a sorting algorithm,” Alison says. He scribbles a list of positive integers in dry erase:

4, 1, 3, 2, 5

The simple task at hand: devise an algorithm to sort these numbers into ascending order. “Human beings know how to do this,” Alison says. “We just kind of do it in our heads.”

Computers, however, must be told precisely how. That requires an algorithm: a set of concrete instructions by which a given problem may be solved. The algorithm Alison shows me is called “bubble sort,” and it works like this:

1. For each number in the set, starting with the first one, compare it to the number that follows, and see if they’re in the desired order.
2. If not, reverse them.
3. Repeat steps 1 and 2 until you’re able to proceed through the set from start to end without reversing any numbers.

The virtue of bubble sort is its simplicity. The downside: If your data set is large, it’s computationally inefficient and time-consuming. Facebook, for obvious reasons, does not use bubble sort. It does use a sorting algorithm to order the set of all posts that could appear in your news feed when you open the app. But that’s the trivial part—a minor subalgorithm within the master algorithm. The nontrivial part is assigning all those posts a numerical value in the first place. That, in short, is the job of the news feed ranking team: to devise a system capable of assigning any given Facebook post a “relevancy score” specific to any given Facebook user.

That’s a hard problem, because what’s relevant to you—a post from your childhood friend or from a celebrity you follow—might be utterly irrelevant to me. For that, Alison explains, Facebook uses a different kind of algorithm, called a prediction algorithm. (Facebook’s news feed algorithm, like Google’s search algorithm or Netflix’s recommendation algorithm, is really a sprawling complex of software made up of smaller algorithms.)

“Let’s say I ask you to pick the winner of a future basketball game, Bulls vs. Lakers,” Alison begins. “Bulls,” I blurt. Alison laughs, but then he nods vigorously. My brain has taken his input and produced an immediate verbal output, perhaps according to some impish algorithm of its own. (The human mind’s algorithms are far more sophisticated than anything Silicon Valley has yet devised, but they’re also **heavily reliant on heuristics** and notoriously prone to folly.)

Random guessing is fine when you’ve got nothing to lose, Alison says. But let’s say there was a lot of money riding on my basketball predictions, and I was making them millions of times a day. I’d need a more systematic approach. “You’re probably going to start by looking at historical data,” he says. “You’re going to look at the win-loss record of each team, the records of the individual players, who’s injured, who’s on a streak.” Maybe you’ll take into account environmental factors: Who’s the home team? Is one squad playing on short rest, or after a cross-country flight? Your prediction algorithm might incorporate all of these factors and more. If it’s good, it will not only predict the game’s winner, but tell you its degree of confidence in the result.

That’s analogous to what Facebook’s news feed algorithm does when it tries to predict whether you’ll like a given post. I ask Alison how many variables—“features,” in machine-learning lingo—Facebook’s algorithm takes into account. “Hundreds,” he says.

It doesn’t just predict whether you’ll actually hit the like button on a post based on your past behavior. It also predicts whether you’ll click, comment, share, or hide it, or even mark it as spam. It will predict each of these outcomes, and others, with a certain degree of confidence, then combine them all to produce a single relevancy score that’s specific to both you and that post. Once every possible post in your feed has received its relevancy score, the sorting algorithm can put them in the order that you’ll see them on the screen. The post you see at the top of your feed, then, has been chosen over thousands of others as the one most likely to make you laugh, cry, smile, click, like, share, or comment.

Yet no matter how meticulously you construct an algorithm, there are always going to be data to which you aren’t privy: the coaches’ game plans, how **Derrick Rose’s knee** is feeling that day, whether **the ball is properly inflated**. In short, the game isn’t played by data. It’s played by people. And people are too complex for any algorithm to model.

Facebook’s prediction algorithm faces still another complication, this one a little more epistemological. The relevancy score is meant to be analogous to the likelihood that the Bulls will win the game. That’s a discrete outcome that’s fully measurable: They either win or they don’t. Facebook’s ranking algorithm used to try to predict a similarly measurable outcome: whether you’d interact in some way with the post in question. Interactions, the humans behind Facebook’s news feed figured, are a good indicator that a given post has struck a chord. They also happen to be the fuel that drives the Facebook economy: clicks, likes, shares, and comments are what make posts go viral, turn individual users into communities, and drive traffic to the advertisers that Facebook relies on for revenue.

But those interactions are only a rough proxy for what Facebook users actually want. What if people “like” posts that they don’t really like, or click on stories that turn out to be unsatisfying? The result could be a news feed that optimizes for virality, rather than quality—one that feeds users a steady diet of candy, leaving them dizzy and a little nauseated, liking things left and right but gradually growing to hate the whole silly game. How do you optimize against that?



It was late 2013, and Facebook was the hottest company in the world. The social network had blown past 1 billion users and gone public at a valuation of more than \$100 billion. It had spent the past year building a revamped mobile app that quickly surpassed Google Search and Google Maps as the nation's most popular. No longer just a way to keep in touch with friends, Facebook had become, in effect, the global newspaper of the 21st century: an up-to-the-minute feed of news, entertainment, and personal updates from friends and loved ones, automatically tailored to the specific interests of each individual user.

Inside the company, the people in charge of the news feed were thrilled with the growth. But while users' engagement was skyrocketing, it wasn't clear whether their overall satisfaction with Facebook was keeping pace. People were liking more things on Facebook than ever. But were they liking Facebook less?

To understand how that question arose, you have to rewind to 2006. Facebook—which was originally little more than a massive compendium of profile pages and groups, something like Myspace—built the news feed in that year as a hub for updates about your friends' activities on the site. Users bristled at the idea that their status updates, profile picture changes, and flirtatious notes on one another's pages would be blasted into the feeds of all of their friends, but Facebook pressed on.

Even then, not everything your friends did made it into your news feed. To avoid overwhelming people with hundreds of updates every day, Facebook built a crude algorithm to filter them based on how likely they were to be of interest. With no real way to measure that—the like button came three years later—the company's engineers simply made assumptions based on their own intuition. Early criteria for inclusion of a post in your news feed included how recent it was and how many of your friends it mentioned. Over time, the team tried tweaking those assumptions and testing how the changes affected the amount of time users spent on the site. But with no way to assess which sorts of posts were delighting people and which were boring, offending, or confusing them, the engineers were essentially throwing darts.



Adam Mosseri (standing), the director of product for news feed, with his team of product managers. Clockwise from far left: Greg Marra, Vibhi Kant, Uzma Barlaskar, Chris Tosswill, Mosseri, Scott Thomson, and Sammi Krug.

Photo by Christophe Wu/Facebook

The like button wasn't just a new way for users to interact on the site. It was a way for Facebook to enlist its users in solving the problem of how best to filter their own news feeds. That users didn't realize they were doing this was perhaps the most ingenious part. If Facebook had told users they had to rank and review their friends' posts to help the company determine how many other people should see them, we would have found the process tedious and distracting. Facebook's news feed algorithm was one of the first to surreptitiously enlist users in personalizing their experience—and influencing everyone else's.

Suddenly the algorithm had a way to identify the most popular posts—and make them go “viral,” a term previously applied to things that were communicated from person to person, rather than broadcast algorithmically to a mass audience. Yet Facebook employees weren’t the only ones who could see what it took for a given post to go viral. Publishers, advertisers, hoaxsters, and even individual users began to glean the elements that viral posts tended to have in common—the features that seemed to trigger reflexive likes from large numbers of friends, followers, and even random strangers. Many began to tailor their posts to get as many likes as possible. Social-media consultants sprung up to advise people on how to game Facebook’s algorithm: the right words to use, the right time to post, the right blend of words and pictures. “LIKE THIS,” a feel-good post would implore, and people would do it, even if they didn’t really care that much about the post. It wasn’t long before Facebook users’ feeds **began to feel eerily similar**: all filled with content that was **engineered to go viral**, much of it mawkish or patronizing. Drowned out were substance, nuance, sadness, and anything that provoked thought or emotions beyond a simple thumbs-up.

Engagement metrics were up—way up—but was this really what the news feed should be optimizing for? The question preoccupied Chris Cox, an early Facebook employee and the news feed’s intellectual architect. “Looking at likes, clicks, comments, and shares is one way of determining what people are interested in,” Cox, 33, tells me via email. (He’s now Facebook’s chief product officer.) “But we knew there were places where this was imperfect. For example, you may read a tragic post that you don’t want to click like, comment on, or share, but if we asked you, you would say that it really mattered to you to have read it. A couple of years ago, we knew we needed to look at more than just likes and clicks to improve how News Feed worked for these kinds of cases.”

An algorithm can optimize for a given outcome, but it can’t tell you what that outcome should be. Only humans can do that. Cox and the other humans behind Facebook’s news feed decided that their ultimate goal would be to show people all the posts that really matter to them and none of the ones that don’t. They knew that might mean sacrificing some short-term engagement—and maybe revenue—in the name of user satisfaction. With Facebook raking in money, and **founder and CEO Mark Zuckerberg controlling a majority of the voting shares**, the company had the rare luxury to optimize for long-term value. But that still left the question of how exactly to do it.

Media organizations have historically defined what matters to their audience through their own editorial judgment. Press them on what makes a story worthwhile, and they’ll appeal to values such as truth, newsworthiness, and public interest. But Cox and his colleagues at Facebook have taken pains to avoid putting their own editorial stamp on the news feed. Instead, their working definition of what matters to any given Facebook user is just this: what he or she would rank at the top of their feeds given the choice. “The perfect way to solve this problem would be to ask everyone which stories they wanted to see and which they didn’t, but that’s not possible or practical,” Cox says. Instead, Facebook decided to ask *some* people which stories they wanted to see and which they didn’t. There were about 1,000 of those people, and until recently, most of them lived in Knoxville, Tennessee. Now they’re everywhere.



Adam Mosseri, Facebook’s 32-year-old director of product for news feed, is Alison’s less technical counterpart—a “fuzzie” rather than a “techie,” in Silicon Valley parlance. He traffics in problems and generalities, where Alison deals in solutions and specifics. He’s the news feed’s resident philosopher.

The push to humanize the news feed’s inputs and outputs began under Mosseri’s predecessor, Will Cathcart. (I wrote about **several of those innovations here**.) Cathcart started by gathering more subtle forms of behavioral data: not just whether someone clicked, but how long he spent reading a story once he clicked on it; not just whether he liked it, but whether he liked it before or after reading. For instance: Liking a post before you’ve read it, Facebook learned, corresponds much more weakly to your actual sentiment than liking it afterward.

After taking the reins in late 2013, Mosseri’s big initiative was to set up what Facebook calls its “feed quality panel.” It began in summer 2014 as a group of several hundred people in Knoxville whom the company paid to come in to an office every day and provide continual, detailed feedback on what they saw in their news feeds. (Their location was, Facebook says, a “historical accident” that grew out of a pilot project in which the company partnered with an unnamed third-party subcontractor.) Mosseri and his team didn’t just study their

behavior. They also asked them questions to try to get at *why* they liked or didn't like a given post, *how much* they liked it, and what they would have preferred to see instead. "They actually write a little paragraph about every story in their news feed," notes Greg Marra, product manager for the news feed ranking team. (This is the group that's becoming Facebook's equivalent of Nielsen families.)

"The question was, 'What might we be missing?'" Mosseri says. "'Do we have any blind spots?'" For instance, he adds, "We know there are some things you see in your feed that you loved and you were excited about, but you didn't actually interact with." Without a way to measure that, the algorithm would devalue such posts in favor of others that lend themselves more naturally to likes and clicks. But what signal could Facebook use to capture that information?



Max Eulenstein, product manager, and Lauren Scissors, user experience researcher, in a conference room near their team's work space.

Photo by Christophe Wu/Facebook

Mosseri deputized product manager Max Eulenstein and user experience researcher Lauren Scissors to oversee the feed quality panel and ask it just those sorts of questions. For instance, Eulenstein used the panel to test the hypothesis that the time a user spends looking at a story in her news feed might be a good indicator that she likes it, even if she didn't actually click like. "We speculated that it might be, but you could think of reasons why it wouldn't be, too," Eulenstein tells me. "It might be that there are scary or shocking stories that you stare at, but don't want to see." The feed quality panelists' ratings allowed Eulenstein and Scissors to not only confirm their hunch, but to examine the subtleties in the correlation, and to begin to quantify it. "It's not as simple as, '5 seconds is good, 2 seconds is bad,'" Eulenstein explains. "It has more to do with the amount of time you spend on a story relative to the other stories in your news feed." The research also revealed the need to control for the speed of users' Internet connections, which can make it seem like they're spending a long time on a given story when they're actually just waiting for the page to load. Out of that research emerged **a tweak that Facebook revealed in June**, in which the algorithm boosted the rankings of stories that users spent more time viewing in their feeds.

Within months, Mosseri and his team had grown so reliant on the panel's feedback that they took it nationwide, paying a demographically representative sample of people around the country to rate and review their Facebook feeds on a daily basis from their own homes. By late summer 2015, Facebook disbanded the Knoxville group and began to expand the feed quality panel overseas. Mosseri's instinct was right: The news feed algorithm had blind spots that Facebook's data scientists couldn't have identified on their own. It took a different kind of data—qualitative human feedback—to begin to fill them in.



Crucial as the feed quality panel has become to Facebook's algorithm, the company has grown increasingly aware that no single source of data can tell it everything. It has responded by developing a sort of checks-and-balances system in which every news feed tweak must undergo a battery of tests among different types of audiences, and be judged on a variety of different metrics.

That balancing act is the task of the small team of news feed ranking engineers, data scientists, and product managers who come to work every day in Menlo Park. They're people like Sami Tas, a software engineer whose job is to translate the news feed ranking team's proposed changes into language that a computer can understand. This afternoon, as I look over his shoulder, he's walking me through a problem that might seem so small as to be trivial. It is exactly the sort of small problem, however, that Facebook now considers critical.

Most of the time, when people see a story they don't care about in their news feed, they scroll right past it. Some stories irk them enough that they're moved to click on the little drop-down menu at the top right of the post and select "Hide post." Facebook's algorithm considers that a strong negative signal and endeavors to show them fewer posts like that in the future.

Not everyone uses Facebook the same way, however. Facebook's data scientists were aware that a small proportion of users—5 percent—were doing 85 percent of the hiding. When Facebook dug deeper, it found that a small subset of those 5 percent were hiding almost every story they saw—even ones they had liked and commented on. For these "superhiders," it turned out, hiding a story didn't mean they disliked it; it was simply their way of marking the post "read," like archiving a message in Gmail.

Yet their actions were biasing the data that Facebook relied on to rank stories. Intricate as it is, the news feed algorithm does not attempt to individually model each user's behavior. It treats your likes as identical in value to mine, and the same is true of our hides. For the superhiders, however, the ranking team decided to make an exception. Tas was tasked with tweaking the code to identify this small group of people and to discount the negative value of their hides.



The news feed engineering and design team. Left to right, back row: Shilin Ding, Meihong Wang, David Vickrey, Lars Backstrom, and Sami Tas. Left to right, front row: Amey Dharwadker, Geoff Teehan, and Sanjeet Hajarnis.

That might sound like a simple fix. But the algorithm is so precious to Facebook that every tweak to the code must be tested—first in an offline simulation, then among a tiny group of Facebook employees, then on a small fraction of all Facebook users—before it goes live. At each step, the company collects data on the change’s effect on metrics ranging from user engagement to time spent on the site to ad revenue to page-load time. Diagnostic tools are set up to detect an abnormally large change on any one of these crucial metrics in real time, setting off a sort of internal alarm that automatically notifies key members of the news feed team.

Once a change like Tas’ has been tested on each of these audiences, he’ll present the resulting data at one of the news feed team’s weekly “ranking meetings” and field a volley questions from Mosseri, Allison, Marra, and his other colleagues as to its effect on various metrics. If the team is satisfied that the change is a positive one, free of unintended consequences, the engineers in charge of the code on the iOS, Android, and Web teams will gradually roll it out to the public at large.

Even then, Facebook can’t be sure that the change won’t have some subtle, longer-term effect that it had failed to anticipate. To guard against this, it maintains a “holdout group”—a small proportion of users who don’t see the change for weeks or months after the rest of us.

To speak of Facebook’s news feed algorithm in the singular, then, can be misleading. It isn’t just that the algorithm is really a collection of hundreds of smaller algorithms solving the smaller problems that make up the larger problem of what stories to show people. It’s that, thanks to all the tests and holdout groups, there are more than a dozen different versions of that master algorithm running in the world at any given time. Tas’ “hide stories” tweak was announced July 31, and his post about it on **Facebook’s “News Feed FYI” blog** passed largely unnoticed by the public at large. Presumably, however, the superhidiers of the world are now marginally more satisfied with their news feeds, and thus more likely to keep using Facebook, sharing stories with friends, and viewing the ads that keep the company in business.



Facebook’s feed quality panel has given the company’s news feed team richer, more human data than it ever had before. Tas and the rest of the ranking team are growing more skillful at finding and fixing the algorithm’s blind spots. But there is one other group of humans that Facebook is turning to more and more as it tries to keep the news feed relevant: ordinary users like you and me.

The survey that Facebook has been running over the past six months—asking a subset of users to choose their favorite among two side-by-side posts—is an attempt to gather the same sort of data from a much wider sample than is possible through the feed quality panel. But the increasing involvement of ordinary users isn’t only on the input side of the equation. Over the past two years, Facebook has been giving users more power to control their news feeds’ output as well.

The algorithm is still the driving force behind the ranking of posts in your feed. But Facebook is increasingly giving users the ability to fine-tune their own feeds—a level of control it had long resisted as onerous and unnecessary. Facebook has spent seven years working on improving its ranking algorithm, Mosseri says. It has machine-learning wizards developing logistic regressions to interpret how users’ past behavior predicts what posts they’re likely to engage with in the future. “We could spend 10 more years—and we will—trying to improve those [machine-learning techniques],” Mosseri says. “But you can get a lot of value right now just by simply asking someone: ‘What do you want to see? What do you not want to see? Which friends do you always want to see at the top of your feed?’ ”

Those are now questions that Facebook allows every user to answer for herself. You can now “unfollow” a friend whose posts you no longer want to see, “see less” of a certain kind of story, and designate your favorite friends and pages as “see first,” so that their posts will appear at the top of your feed every time you log in. How to do all of these things is not immediately obvious to the casual user: You have to click a tiny gray down arrow in the top right corner of a post to see those options. Most people never do. But as the limitations of the fully automated feed have grown clearer, Facebook has grown more comfortable highlighting these options via occasional pop-up reminders with links to explanations and help pages. It is also testing new ways for users to interact with the news feed, including alternate, **topic-based news feeds** and new buttons to convey **reactions other than like**.

The shift is partly a defensive one. The greatest challenges to Facebook’s dominance in recent years—the upstarts that threaten to do to Facebook what Facebook did to Myspace—have eschewed this sort of data-driven approach altogether. Instagram, which Facebook acquired in 2012 in part to **quell the threat** posed by its fast-growing popularity, simply shows you every photo from every person you follow in chronological order. Snapchat has eclipsed Facebook as teens’ social network of choice by eschewing virality and automated filtering in favor of more intimate forms of digital interaction.

Facebook is not the only data-driven company to run up against the limits of algorithmic optimization in recent years. Netflix’s famous movie-recommendation engine has come to rely heavily on humans who are paid to watch movies all day and classify them by genre. To counterbalance the influence of Amazon’s automated A/B tests, CEO Jeff Bezos places outsized importance on the specific complaints of individual users and maintains **a public email address** for that very purpose. It would be premature to declare the age of the algorithm over before it really began, but there has been a change in velocity. Facebook’s Mosseri, for his part, rejects the buzzword “data-driven” in reference to decision making; he prefers “data-informed.”



Engineering director Tom Alison, product director Adam Mosseri, and engineering director Lars Backstrom in the Saint Frank coffee shop in Facebook’s Building 20.

Photo by Christophe Wu/Facebook

Facebook’s news feed ranking team believes the change in its approach is paying off. “As we continue to improve news feed based on what people tell us, we are seeing that we’re getting better at ranking people’s news feeds; our ranking is getting closer to how people would rank stories in their feeds themselves,” says Scissors, the user experience researcher who helps to oversee the feed quality panel.

There’s a potential downside, however, to giving users this sort of control: What if they’re mistaken, as humans often are, about what they really want to see? What if Facebook’s database of our online behaviors really did know us better, at least in some ways, than we knew ourselves? Could giving people the news feed they say they want actually make it less addictive than it was before?

Mosseri tells me he’s not particularly worried about that. The data so far, he explains, suggest that placing more weight on surveys and giving users more options have led to an increase in overall engagement and time spent on the site. While the two goals may seem to be in tension in the short term, “We find that qualitative improvements to the news feed look like they correlate with long-term engagement.” That may be a happy coincidence if it continues to hold true. But if there’s one thing that Facebook has learned in 10 years of running the news feed, it’s that data never tell the full story, and the algorithm will never be perfect. What looks like it’s working today might be unmasked as a mistake tomorrow. And when it does, the humans who go to work every day in Menlo Park will read a bunch of spreadsheets, hold a bunch of meetings, run a bunch of tests—and then change the algorithm once again.