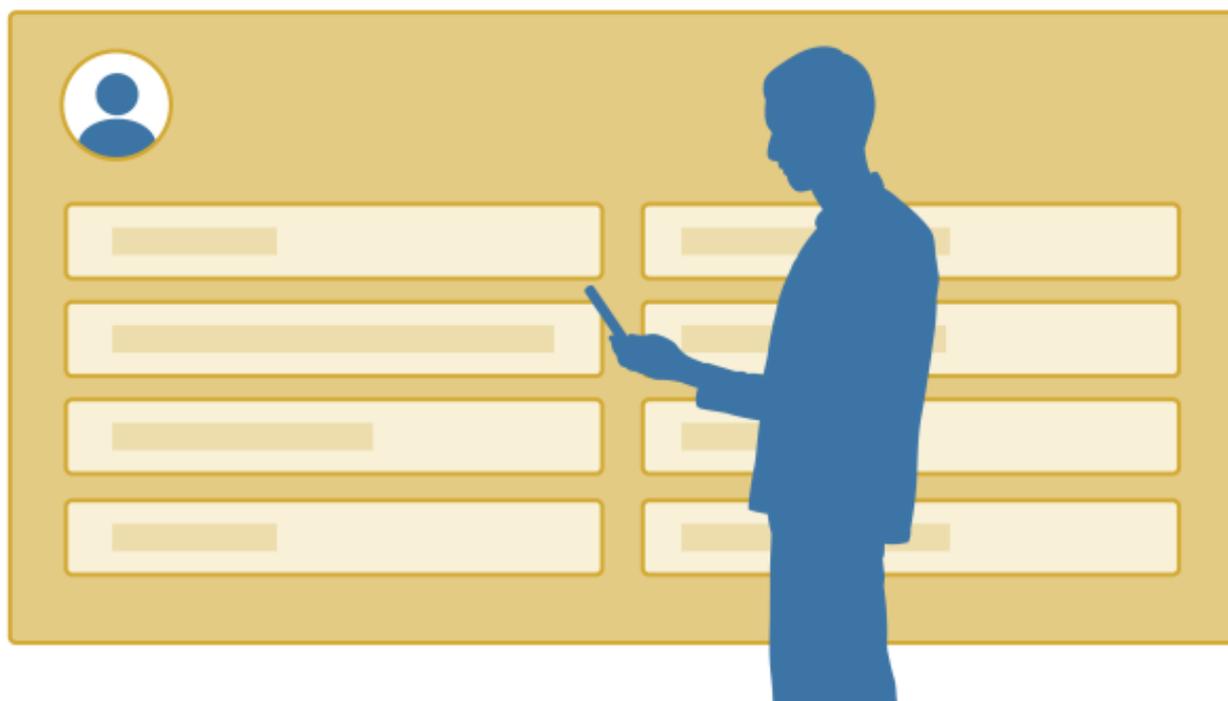


Facebook Algorithms and Personal Data

About half of Facebook users say they are not comfortable when they see how the platform categorizes them, and 27% maintain the site's classifications do not accurately represent them

By [Paul Hitlin](#) and [Lee Rainie](#)



Most commercial sites, from social media platforms to news outlets to online retailers, collect a wide variety of data about their users' behaviors. Platforms use this data to deliver content and recommendations based on users' interests and traits, and to allow advertisers to [target ads](#) to relatively precise segments of the public. But how well do Americans understand these algorithm-driven classification systems, and how much do they think their lives line up with what gets reported about them? As a window into this hard-to-study phenomenon, a new Pew Research Center

survey asked a representative sample of users of the nation’s most popular social media platform, Facebook, to reflect on the data that had been collected about them. (See more about why we study Facebook in the box below.)

Facebook makes it relatively easy for users to find out how the site’s algorithm has categorized their interests via a “[Your ad preferences](#)” page.¹ Overall, however, 74% of Facebook users say they did not know that this list of their traits and interests existed until they were directed to their page as part of this study.

When directed to the “ad preferences” page, the large majority of Facebook users (88%) found that the site had generated some material for them. A majority of users (59%) say these categories reflect their real-life interests, while 27% say they are not very or not at all accurate in describing them. And once shown how the platform classifies their interests, roughly half of Facebook users (51%) say they are not comfortable that the company created such a list.

The survey also asked targeted questions about two of the specific listings that are part of Facebook’s classification system: users’ political leanings, and their racial and ethnic “affinities.”

In both cases, more Facebook users say the site’s categorization of them is accurate than say it is inaccurate. At the same time, the findings show that portions of users think Facebook’s listings for them are not on the mark.

Many Facebook users say they do not know the platform classifies their interests, and roughly half are not comfortable with being categorized

% of U.S. adult Facebook users who say ____ after being directed to view their Facebook ‘ad preferences’ page

74% They **did not know** Facebook maintained this list of their interests and traits

51% They are **not comfortable** with Facebook compiling this information

27% The listings **do not very or at all accurately represent** them

Source: Survey of Facebook users conducted Sept. 4-Oct. 1, 2018. “Facebook Algorithms and Personal Data”

PEW RESEARCH CENTER

When it comes to politics, about half of Facebook users (51%) are assigned a political “affinity” by the site. Among those who are assigned a political category by the site, 73% say the platform’s categorization of their politics is very or somewhat accurate, while 27% say it describes them not very or not at all accurately. Put differently, 37% of Facebook users are both assigned a political affinity *and* say that affinity describes them well, while 14% are both assigned a category and say it does not represent them accurately.

For some users, Facebook also lists a category called “multicultural affinity.” According to third-party online courses about how to [target ads](#) on Facebook, this listing is meant to designate a user’s “affinity” with various racial and ethnic groups, rather than assign them to groups reflecting their *actual* race or ethnic background. Only about a fifth of Facebook users (21%) say they are listed as having a “multicultural affinity.” Overall, 60% of users who are assigned a multicultural affinity category say they do in fact have a very or somewhat strong affinity for the group to which they are assigned, while 37% say their affinity for that group is not particularly strong. Some 57% of those who are assigned to this category say they do in fact consider themselves to be a member of the racial or ethnic group to which Facebook assigned them.

These are among the findings from a survey of a nationally representative sample of 963 U.S. Facebook users ages 18 and older conducted Sept. 4 to Oct. 1, 2018, on GfK’s KnowledgePanel.

A second survey of a representative sample of all U.S. adults who use social media – including Facebook and other platforms like Twitter and Instagram – using Pew Research Center’s [American Trends Panel](#) gives broader context to the insights from the Facebook-specific study.

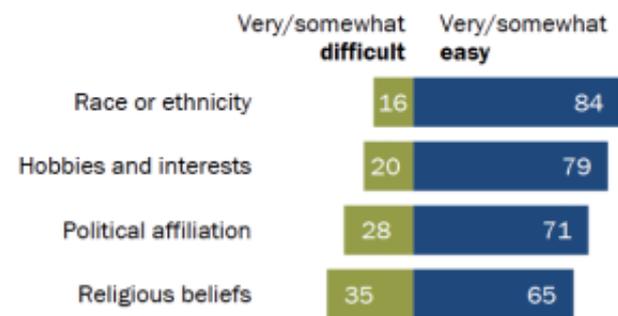
This second survey, conducted May 29 to June 11, 2018, reveals that social media users generally believe it would be relatively easy for the platforms they use to determine key traits about them based on the data they have

amassed about their behaviors.

Majorities of social media users say it would be very or somewhat easy for these platforms to determine their race or ethnicity (84%), their hobbies and interests (79%), their political affiliation (71%) or their religious beliefs (65%). Some 28% of social media users believe it would be difficult for these platforms to figure out their political views, nearly matching the share of Facebook users who are assigned a political listing but believe that listing is not very or not at all accurately.

Social media users say it is easy for sites to identify their race and interests

% of U.S. adult social media users who say it would be ___ for social media sites to figure out their ...



Note: Respondents who did not give an answer are not shown.
Source: Survey conducted May 29-June 11, 2018.
"Facebook Algorithms and Personal Data"

PEW RESEARCH CENTER

Why we study Facebook

Pew Research Center chose to study Facebook for this research on public attitudes about digital tracking systems and algorithms for a number of reasons. For one, the platform is used by a considerably [bigger](#) number of Americans than other [popular social media platforms](#) like Twitter and Instagram. Indeed, its global user base is [bigger than the population](#) of many countries. Facebook is the third most trafficked website in the [world](#) and fourth most in the [United States](#). Along with [Google](#), Facebook dominates the [digital advertising](#) market, and the firm itself [elaborately documents](#) how advertisers can micro-target audience segments. In addition, the Center's studies have shown that Facebook holds a special and meaningful place in the [social](#) and [civic](#) universe of its users.

The company allows users to view at least a partial compilation of how it classifies them on the page called "[Your ad preferences](#)." It is relatively simple to find this page, which allows researchers to direct Facebook users

to their preferences page and ask them about what they see.

Users can find their own preferences page by following the directions in the [Methodology section](#) of this report. They can [opt out](#) of being categorized this way for ad targeting, but they will still get other kinds of [less-targeted ads](#) on Facebook.

Most Facebook users say they are assigned categories on their ad preferences page

A substantial share of websites and apps [track](#) how people use digital services, and they use that data to deliver services, content or advertising targeted to those with specific interests or traits. Typically, the precise workings of the proprietary algorithms that perform these analyses are unknowable outside the companies who use them. At the same time, it is clear the process of algorithmically assessing users and their interests involves a lot of informed guesswork about the meaning of a user's activities and how those activities add up to elements of a user's identity.

Facebook, the most prominent social network in the world, analyzes scores of different dimensions of its [users' lives](#) that advertisers are then invited to target. The company allows users to view at least a partial compilation of how it classifies them on the page called "[Your ad preferences](#)." The page, which is different for each user, displays several types of personal information about the individual user, including "your categories" – a list of a user's purported interests crafted by Facebook's algorithm. The categorization system takes into account data provided by users to the site and their engagement with content on the site, such as the material they have posted, liked, commented on and shared.

These categories might also include insights Facebook has gathered from a user's online behavior outside of the Facebook platform. Millions of [companies and organizations](#) around the world have activated the [Facebook pixel](#) on their websites. The Facebook pixel records the activity of

Facebook users on these websites and passes this data back to Facebook. This information then allows the companies and organizations who have activated the pixel to better target advertising to their website users who also use the Facebook platform. Beyond that, Facebook has [a tool](#) allowing advertisers to link offline conversions and purchases to users – that is, track the offline activity of users after they saw or clicked on a Facebook ad – and find audiences similar to people who have converted offline. (Users can opt out of having their information used by this targeting feature.)

Overall, the array of information can cover users' demographics, social networks and relationships, political leanings, life events, food preferences, hobbies, entertainment interests and the digital devices they use.

[Advertisers](#) can select from these categories to target groups of users for their messages. The existence of this material on the Facebook profile for each user allows researchers to work with Facebook users to explore their own digital portrait as constructed by Facebook.

The Center's representative sample of American Facebook users finds that 88% say they are assigned categories in this system, while 11% say that after they are directed to their ad preferences page they get a message saying, "You have no behaviors."

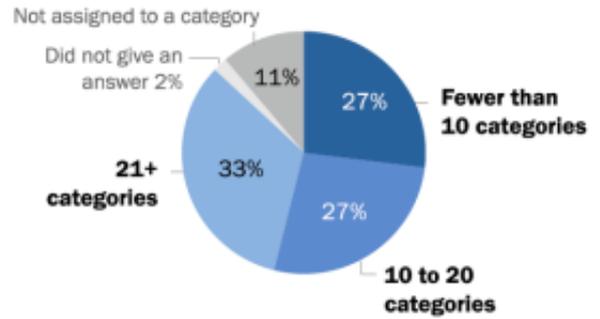
Some six-in-ten Facebook users report their preferences page lists either 10 to 20 (27%) or 21 or more (33%) categories for them, while 27% note their list contains fewer than 10 categories.

Those who are heavier users of Facebook and those who have used the site the longest are more likely to be listed in a larger number of personal interest categories. Some 40% of those who use the platform multiple times a day are listed in 21 or more categories, compared with 16% of those who are less-than-daily users. Similarly, those who have been using Facebook for 10 years or longer are more than twice as likely as those with less than five years of experience to be listed in 21 or more categories (48% vs. 22%).

74% of Facebook users say they did not know about the platform's list of their interests

A majority of Facebook users have 10 or more categories listed on their ad preferences page

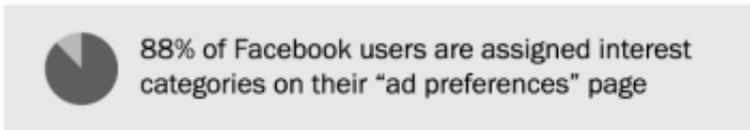
% of U.S. adult Facebook users with ____ categories listed on their 'ad preferences' page



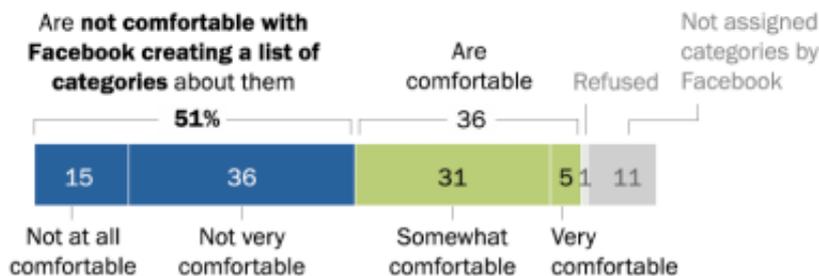
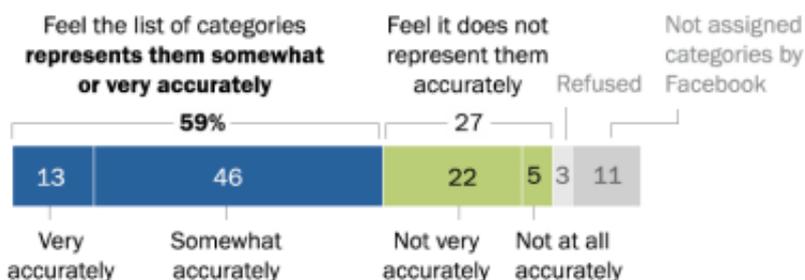
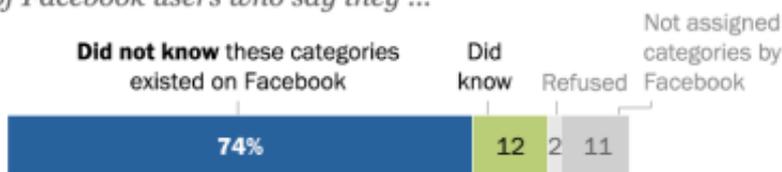
Source: Survey of Facebook users conducted Sept. 4-Oct. 1, 2018. "Facebook Algorithms and Personal Data"

PEW RESEARCH CENTER

Most Facebook users do not know the platform lists their interests for advertisers, and half are not comfortable with these lists



% of Facebook users who say they ...



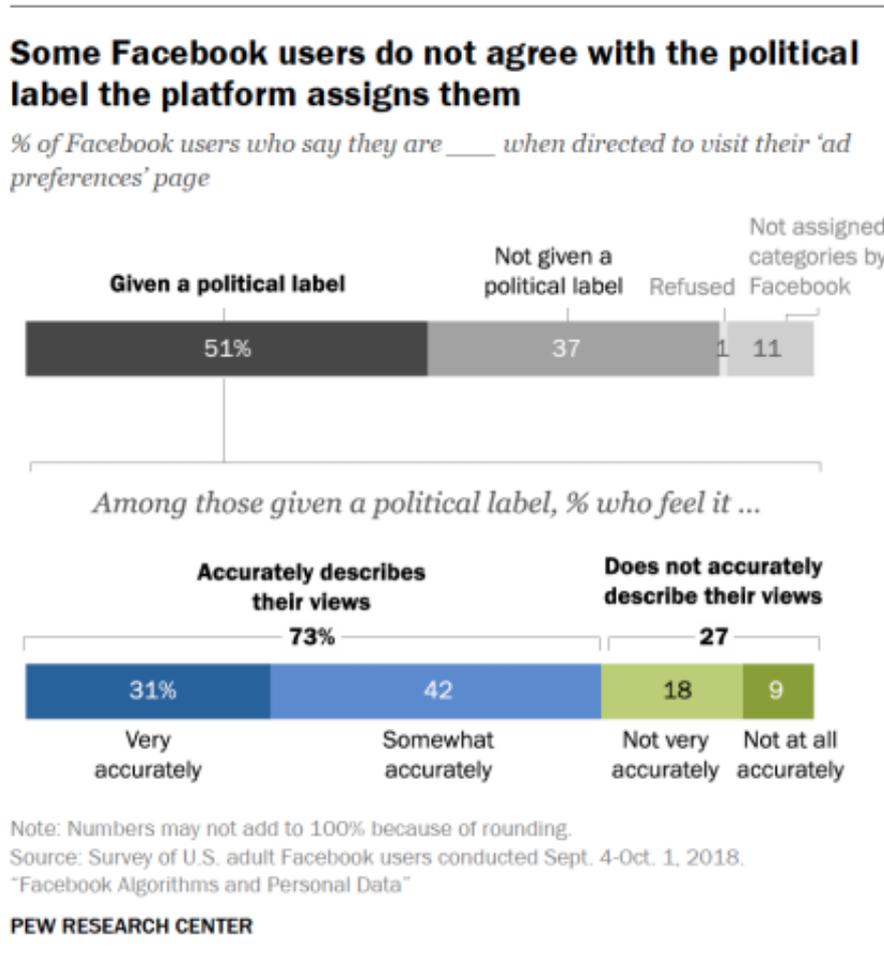
Note: Numbers may not add to 100% because of rounding.
 Source: Survey of U.S. adult Facebook users conducted Sept. 4-Oct. 1, 2018.
 "Facebook Algorithms and Personal Data"

PEW RESEARCH CENTER

About three-quarters of Facebook users (74%) say they did not know this list of categories existed on Facebook before being directed to the page in the Center’s survey, while 12% say they were aware of it.² Put differently, 84% of those who reported that Facebook had categorized their interests did not know about it until they were directed to their ad preferences page.

When asked how accurately they feel the list represents them and their

interests, 59% of Facebook users say the list very (13%) or somewhat (46%) accurately reflects their interests. Meanwhile, 27% of Facebook users say the list not very (22%) or not at all accurately (5%) represents them.



Yet even with a majority of users noting that Facebook at least somewhat accurately assesses their interests, about half of users (51%) say they are not very or not at all comfortable with Facebook creating this list about their interests and traits. This means that 58% of those whom Facebook categorizes are not generally comfortable with that process. Conversely, 5% of Facebook users say they are very comfortable with the company creating this list and another 31% declare they are somewhat comfortable.

There is clear interplay between users' comfort with the Facebook traits-assignment process and the accuracy they attribute to the process. About three-quarters of those who feel the listings for them are not very or not at

all accurate (78%) say they are uncomfortable with lists being created about them, compared with 48% of those who feel their listing is accurate.

Facebook’s political and ‘racial affinity’ labels do not always match users’ views

It is relatively common for Facebook to assign political labels to its users. Roughly half (51%) of those in this survey are given such a label. Those assigned a political label are roughly equally divided between those classified as liberal or very liberal (34%), conservative or very conservative (35%) and moderate (29%).

Among those who are assigned a label on their political views, close to three-quarters (73%) say the listing very accurately or somewhat accurately describes their views. Meanwhile, 27% of those given political classifications by Facebook say that label is not very or not at all accurate.

There is some variance between what users say about their political ideology and what Facebook attributes to them.³ Specifically, self-described moderate Facebook users are more likely than others to say they are not classified accurately. Among those assigned a political category, some 20% of self-described liberals and 25% of those who describe themselves as conservative say they are not described well by the labels Facebook assigns to them. But that share rises to 36% among self-described moderates.

In addition to categorizing users’ political views, Facebook’s algorithm assigns some users to groups by “multicultural affinity,” which the firm says it assigns to people whose Facebook activity “aligns with” certain cultures. About one-in-five Facebook users (21%) say they are assigned such an affinity.

The use of multicultural affinity as a tool for advertisers to exclude certain groups has created controversies. Following pressure from [Congress](#) and

investigations by [ProPublica](#), Facebook [signed an agreement](#) in July 2018 with the Washington State Attorney General saying it would no longer let advertisers unlawfully exclude users by race, religion, sexual orientation and other protected classes.

In this survey, 43% of those given an affinity designation are said by Facebook's algorithm to have an interest in African American culture, and the same share (43%) is assigned an affinity with Hispanic culture. One-in-ten are assigned an affinity with Asian American culture. Facebook's [detailed targeting tool](#) for ads does not offer affinity classifications for any other cultures in the U.S., including Caucasian or white culture.

Of those assigned a multicultural affinity, 60% say they have a “very” or “somewhat” strong affinity for the group they were assigned, compared with 37% who say they do not have a strong affinity or interest.⁴ And 57% of those assigned a group say they consider themselves to be a member of that group, while 39% say they are not members of that group.