Lyra Health Ex-Therapists Warn Of Ethical Conflicts
Richard Nieva
February 10, 2022

Feeling stressed and overwhelmed last January, Daniel Rojas decided to take advantage of a benefit Starbucks often touts for its employees around the country: free therapy through Lyra Health, a mental health startup that provides counseling services for some of the biggest companies in the world.

Rojas, a 25-year-old shift supervisor in Buffalo, New York, had been dealing with gender dysphoria and body image problems, two issues he says compound each other "like a snake eating its tail." So Rojas jumped at the coffee giant's offer of 20 free counseling sessions from Lyra, a Silicon Valley darling cofounded and led by former Facebook CFO David Ebersman.

But four sessions in, Rojas, who uses he/they pronouns, felt frustrated with the progress of their treatment. He said he had to constantly re-explain things he'd gone over in previous sessions, which made him relive the same traumas every time he had to repeat them. So they decided to end treatment with that counselor and find another one on Lyra's platform.

When they attempted to find someone else, though, they said a Lyra rep told them in a video call that their issues were too advanced for the company's care. The rep suggested he seek long-term treatment elsewhere and left him to figure it out on his own.

"I work really hard at Starbucks and I want to get every benefit I possibly can," Rojas said. "I felt alienated. I felt like I was being cheated."

Starbucks did not respond to multiple requests for comment on Rojas's situation, and Lyra declined to address it.

The tech industry's growth-at-all-costs outlook may not translate well to a field as delicate as mental health.

Starbucks bills its Lyra benefit as "mental healthcare for a wide-range of needs, from mild to complex." But Rojas's experience reveals one way patients can feel underserved by a startup aiming to be a model for "modern mental healthcare." In interviews with BuzzFeed News, 18 users, therapists, and former Lyra employees voiced concerns about some of the company's business practices, including its productivity-based bonus structure for therapists and its use of patient data. Some of the people who spoke to BuzzFeed News for this story did so under the condition of anonymity because they feared repercussions from their employers or former employers.

Lyra — whose juggernaut slate of corporate clients also includes Google, Facebook parent Meta, and Morgan Stanley — is one of the leaders in a wave of startups focusing on mental health, applying Silicon Valley's data-absorbed ethos to the discipline of therapy. Tech giants like Facebook and Google often weather criticism for taking liberties with people's personal information, but the business model behind startups such as Lyra has received less scrutiny. The company, which has raised $700 million in funding to date, generates revenue through deals with high-profile companies, using anonymized patient data to prove it provides worthwhile benefits.

Better access to therapy, of course, is a good thing. Lyra's supporters cite good wages for therapists, a well-built software platform, and the awareness the company has brought to people who might not have otherwise sought therapy. Other mental health companies, including Ginger, Modern Health, and Cerebral, have also become workplace staples, especially throughout a global pandemic. (BuzzFeed has a relationship with Ginger to offer mental health benefits to employees.)

As more people entrust this burgeoning class of therapy apps with their well-being, the tech industry's growth-at-all-costs outlook may not translate well to a field as delicate as mental health. Lyra's prominence raises questions about whether a high-flying Silicon Valley startup's need to justify its reported $4.6 billion valuation conflicts with its ability to provide quality mental health services.

Lyra spokesperson Dyani Vanderhorst said in a statement, "Our approach makes it easy for millions of people to access high-quality mental healthcare. As demand accelerates, we remain committed to delivering clinically proven, outcomes-based mental healthcare for employees and their dependents across all facets of mental health."

Ebersman founded Lyra Health seven years ago in Burlingame, California, about 20 miles south of San Francisco. The former Facebook executive, who was previously the financial chief at Genentech before arriving at Mark Zuckerberg's social network, said he decided to start Lyra after having a difficult experience finding care for a family member. (Lyra declined to make Ebersman available for an interview.)

The startup employs its own therapists while also tapping into a network of contractors. When a company hires Lyra to be an Employee Assistance Program (EAP), its employees are typically given a set number of free sessions per year to see a counselor. The original plan was to offer users unlimited therapy sessions, two former early employees said, though that policy was later changed. The clinicians on Lyra's platform specialize in evidence-based "blended care" therapy, a mix of in-person or live-video sessions and digital lessons and other content. After employees use all of their free sessions, they can continue seeing their Lyra therapist by paying out of pocket or through health insurance.

When it comes to clinical work, the company puts an emphasis on efficiency. The startup's in-house therapists are entitled to bonuses based on productivity, two former Lyra staff therapists told BuzzFeed News, which is measured through a range of goals, including symptoms improving over time based on patient surveys.

"You can't just throw people in and expect them to see results."

One of the former therapists, Megha Reddy, said the bonus model can push therapists into "churning out" patients quickly. Reddy, who worked at Lyra until 2019, said the system can encourage questionable behavior, and could incentivize therapists to not see a patient for more than a certain number of sessions.

"This isn't an assembly line. This is actually people," Reddy said. "You can't just throw people in and expect them to see results."

Vanderhorst, the Lyra spokesperson, didn't answer specific questions about the bonus system or what changes may have been made to it, but said in a statement, "We take great care in

creating a supportive and dynamic work experience for our providers as well as offering them fair compensation."

As a part-time employee working 20 hours a week at Lyra, Reddy said she was expected to see 12 to 20 patients a week with the goal of having a whole new slate of patients every six to 10 weeks. The financial incentives create the potential for abuse, she said. Her discomfort with the bonus system was her main reason for leaving Lyra.

"It can get dicey in terms of ethics," Reddy said. "You're not going to dictate to me when a patient is supposed to feel better based on numbers. That's going to be based on the patient and my discretion."

Vanderhorst said providers are the ones that determine how many sessions a patient needs.

Arthur Caplan, head of the Division of Medical Ethics at the NYU Grossman School of Medicine, said a bonus system like the one used by Lyra makes him "nervous." "It could be a conflict of interest," he said. "Turnover as a measure of success is certainly dubious in mental healthcare."

"We set the tone. We basically started an industry."

Facebook, Google, and Morgan Stanley declined to comment on Lyra's bonus structure, and Starbucks did not respond to multiple requests for comment.

Other mental health startups have also reportedly incentivized productivity from therapists. In December, Forbes reported that Cerebral had reclassified salaried therapists as contractors, making access to medical, vision, and dental benefits contingent on meeting quotas. "This was done so that our best and most productive therapists have the opportunity to earn more," CEO Kyle Robertson said in response. Cerebral did not respond to a request for comment.

But while other apps engage in similar practices when it comes to data policies and productivity incentives, Lyra Health bears some of the responsibility because it was a pioneer in the space, two former employees said. "We set the tone," said one of them. "We basically started an industry."

Ebersman has said he wants to bring some of Facebook's data-centric approach to mental health. "One of the things that's so magical about Facebook is how the experience is completely personalized," Ebersman said when Lyra launched in 2015. "And that is generally absent from your experience in healthcare."

To collect data on the progress of treatment, Lyra periodically sends patients "outcomes surveys." The questionnaires inquire, for example, about things like anxiety or irritability over the last two weeks, asking patients to rank their intensity from 0 to 3, according to surveys viewed by BuzzFeed News. The surveys, which use clinically accepted and standardized questions, are optional. But patients may feel compelled to complete them because the automated emails look like they are coming from their therapist.

Clinicians can use the data to help shape their treatment, but there's another reason Lyra pushes the surveys: The company shares aggregated and anonymized data about patient outcomes with employers to illustrate the effectiveness of its services.

In one version of the survey viewed by BuzzFeed News that is hosted on research.net, a disclosure that explains how Lyra shares aggregated and anonymous outcomes data with employers appears on page three of five. Another version of the survey accessed through Google's internal Lyra portal and viewed by BuzzFeed News does not explicitly say that outcomes data will be shared. Instead, it reads: "Your responses are confidential and are not shared with the employer sponsoring your Lyra benefit." Lyra declined to answer questions about how it currently discloses to patients that it shares outcomes data with employers.

Google and Starbucks confirmed they receive data from Lyra in order to judge the service's value to employees. "Google does not access the medical records of people using Lyra Health, and we have no special access," Google spokesperson Jennifer Rodstrom said in a statement. Facebook and Morgan Stanley declined to comment.

"The bottom line is, this is a business. So the bottom line is money."

Outcomes data is so central to Lyra's philosophy that the company's previous name was Outcome Health, according to an internal document viewed by BuzzFeed News. The name was changed to Lyra Health prior to the company's launch.

"The bottom line is, this is a business. So the bottom line is money," said one former Lyra employee who worked on the company's clinical team. "And how can you get money? By data. By saying, 'Look how successful we are. Please invest in us.'"

BuzzFeed News spoke to seven current and former Google, Facebook, and Starbucks employees who saw Lyra therapists and were upset about the sharing of outcomes data. One former Facebook employee, who worked on privacy initiatives at the tech giant, was concerned the data could be abused even if aggregated and anonymized. "I understand that employers want to measure the efficacy of their programs," the former employee said, but it's "completely inappropriate" to share such sensitive data.

Aside from the disclosure on some surveys, Lyra has laid out its data practices in a privacy policy, a more than 5,000-word document that lives at the bottom of its website. The company says the data sharing complies with the Health Insurance Portability and Accountability Act of 1996 (HIPAA), which regulates the use of health information. The company's HIPAA notice, also found at the bottom of its website, says Lyra shares patient data "to support our business operations."

Vanderhorst said new users must acknowledge both the privacy policy and HIPAA notice while setting up their accounts.

Still, some patients had not known about the data sharing. Of the seven current and former Google, Facebook, and Starbucks employees who spoke to BuzzFeed News, all but one of them said they did not know the data from these surveys could be shared with employers in any form. "It's shocking to me," said a former Google employee, who said she didn't remember a data disclosure while filling out the surveys. "I had no idea they were doing that."

Lyra defended how it communicates its privacy practices to patients. "Lyra follows all U.S. regulations regarding privacy," Vanderhorst said in a statement. "Our privacy policy is standard format and provides detailed information about our practices."

Jennifer King, privacy and data policy fellow at the Stanford University Institute for Human-Centered Artificial Intelligence, said there's the legal process of gathering consent, and then there's "the moral question" of making sure people fully understand. The added layer of sharing information with an employer makes it even more problematic. "People tend to feel somewhat better with aggregation, but in the workplace is different," she said.

Lyra isn't the only company in the mental health space facing questions about what it's doing with anonymous user data. Loris.ai, the partner company to the nonprofit Crisis Text Line, is contending with criticism after Politico reported that it uses anonymous but sensitive data drawn from conversations on the text-based suicide hotline for business purposes.

Some Lyra therapists were not aware Lyra shares outcomes data with employers, either. BuzzFeed News interviewed eight current and former Lyra therapists, and six of them said they did not know about the data sharing. The therapists said meaningful consent from patients is crucial, even though their names are not attached to the data.

Some patients and therapists didn't mind the data being shared anonymously, since it might be valuable for a company to know if its workforce is depressed or riddled with anxiety. But one former Lyra therapist says patients should get to choose what they want shared. "They should be able to select whether they're willing for their outcomes to be reported," she said.

Data collection was a key issue for some therapists during the early days of the company, according to three former Lyra employees. They said concerns about data sharing made it difficult to recruit therapists to work with Lyra when the company was getting started. When company leadership was told about those hesitations, they were dismissive of the concerns, the former employees said.

"Lyra has tremendous respect for the clinical knowledge, experience, and expertise of our providers," Vanderhorst said in a statement. "Provider recruitment and retention are essential to the care we provide members and the success of our organization."

The company has also had a history of its clinicians feeling overlooked, two former employees said. While engineering and data teams were valued for their input, people on the clinical team were treated like "second-class citizens," one of the former employees said. That employee said that culture was instilled as Ebersman began to bring in people who used to work at Facebook. Lyra did not address these allegations and Facebook declined to comment.

Chelsey Glasson, a former Google and Facebook employee, has recently sounded the alarm on EAPs like Lyra and the potential conflict of interest that could occur when your employer pays for your therapist. In an October op-ed for Insider, she called for more transparency in the relationship between third-party mental health providers and employers. Glasson, who is suing Google after alleged pregnancy discrimination, had sought session notes from her Lyra therapist as part of the lawsuit. Google then demanded and received the notes as well. After that, Glasson said, her therapist called and indicated she was no longer comfortable seeing her.

Google declined to comment. Glasson's former therapist didn't respond to multiple requests for comment. In Lyra's privacy policy, the company says it can use personal information to "comply with our legal obligations."

“It’s all inappropriate and unethical,” Glasson said of Lyra’s business practices. “People have no idea this is happening.”

Glasson, who is based in Seattle, filed a complaint against her therapist, and the situation is now under investigation by the Washington State Department of Health, according to emails viewed by BuzzFeed News.

“It’s all inappropriate and unethical,” Glasson said of Lyra’s business practices. “People have no idea this is happening.”

After consulting with Glasson, Washington State Sen. Karen Keiser sent a letter in November to the state’s secretary of health about the “potential conflict” between employees and employers that participate in EAPs, according to a copy of the letter viewed by BuzzFeed News. Then, in December, Keiser pre-filed legislation that aims to give workers more rights when it comes to EAPs. The bill, called SB 5564, would prohibit employers from disciplining workers based on their decision to see — or not see — a therapist through an EAP. It would also make it illegal for an employer to obtain individually identifiable information about an employee. A state senate committee discussed the bill at a hearing last month.

“Our huge technology companies don't hold personal privacy with the same regard that I think they should,” Keiser told BuzzFeed News. “They’ve been data mining personal privacy information for years and for their own bottom line. But when they use it on their employees, that’s a whole different thing. It’s really a big brother kind of approach.”

Lyra’s policies have at least some people wary about seeking therapy through their employers. After Glasson’s experience with her therapist was reported by the New York Times in July, some Google workers became less likely to use the EAP services provided by Lyra, said the Alphabet Workers Union, which represents workers for Google and its parent company. Google declined to comment.

“I was surprised when I heard about her story,” said a former Google employee. “It really shed a lot of light on the relationship that the counselor has with the company.”

Cyberattacks top list of 2022 health tech hazards alongside supply chain problems, damaged infusion pumps

Rebecca Torrence
January 20, 2022

Cybersecurity is the most pressing technology issue facing health organizations in the new year, according to the ECRI Institute.

The organization's list of the top 10 health technology hazards for 2022, which sets out to identify potential dangers in the use of medical devices and systems, awarded cyberattacks the top spot, followed by supply chain concerns and damaged infusion pumps.

The rankings, beginning with the highest priority concerns, indicate a need to prioritize stronger, more resilient systems and practices in healthcare as organizations continue to recover from pandemic lows.

Cyberattacks, which can endanger day-to-day business operations as well as patient safety, slid into the top spot as a necessary target for every health organization, according to the nonprofit.

As the report reads, "the question is not whether a given facility will be attacked, but when."

Compromised medical devices and data systems can result in the rescheduling of appointments and surgeries, the diversion of emergency vehicles and shutdowns to care units or even entire organizations.

These attacks pose a particular risk in part because many health organizations have not implemented sufficient security measures to anticipate them.

In an August 2021 survey conducted by CyberMDX in collaboration with Philips, less than 11% of hospital IT executives said cybersecurity was a high-priority investment.

At the same time, 48% of those executives reported a forced or proactive shutdown of their systems in the last six months due to ransomware attacks or queries.

"Responding to these risks requires not only a robust security program to prevent attacks from reaching critical devices and systems, but also a plan for maintaining patient care when they do," ECRI said in the report.

Problems with telehealth came in at No. 5 on the list. As many virtual care programs were rapidly launched in 2020 to respond to demand for virtual services during the pandemic, not all have been optimized for care delivery, according to the organization.

It's no secret that telehealth options have ballooned since the start of the pandemic. A Rock Health survey on digital health consumer adoption saw 80% of respondents reporting their primary care provider offered telemedicine in 2021, while only 44% said they were given that choice before the pandemic.

Virtual care looks to be sticking around for the long haul. Rock Health analysts reported that 73% of telemedicine users said they expected to keep using telehealth at the same rate or at a higher frequency in the future.

Facilities banking on the long-term potential of virtual care should prioritize technologies that are easy to use and that promote accurate data collection and transfer, ECRI said.

Providers should also examine whether a telehealth solution truly allows them to deliver care of equal quality to in-person services as well as whether the use of that solution is appropriate for a given patient's needs.

Here is ECRI's full list of health technology hazards for the new year:

1. Cybersecurity attacks
2. Supply chain shortfalls
3. Damaged infusion pumps
4. Inadequate emergency stockpiles
5. Telehealth workflow and human factors shortcomings
6. Failure to adhere to syringe pump best practices
7. AI-based reconstruction
8. Poor duodenoscope reprocessing ergonomics and workflows
9. Disposable gowns with insufficient barrier protection
10. Wi-Fi dropouts and dead zones

UK NHS to test AI systems for biases in healthcare

Jack Aldane
February 14, 2022

As part of the pilot programme, the NHS will support AI researchers and developers with information obtained through engagement with patients and healthcare professionals. Photo by EVG Kowalievska via Pexels

The UK's National Health Service (NHS) is trialling a programme designed to identify algorithmic biases in systems used to administer healthcare.

Its aim is to use Algorithmic Impact Assessments (AIAs) to address decisions made by artificial intelligence (AI) that risk worsening healthcare outcomes for patients based on their profile and background.

"By allowing us to proactively address risks and biases in systems which will underpin the health and care of the future, we are ensuring that we create a system of healthcare which works for everyone, no matter who you are or where you are from," said Syed Kamall, under-secretary of state for innovation.

The body behind the pilot, NHS AI Lab, commissioned the Ada Lovelace Institute to produce a methodology for using AIAs. The institute has since published a research paper that outlines its methodology and aims to help developers and researchers understand the ways in which AI technologies can impact people, society and the environment.

Octavia Reeve, interim lead at the Ada Lovelace Institute, commented: "[These] assessments have the potential to create greater accountability for the design and deployment of AI systems in healthcare, which can in turn build public trust in the use of these systems, mitigate risks of harm to people and groups, and maximise their potential for benefit."

The Ada Lovelace Institute report cites the only model of AIA currently in use, which was authorised by the Treasury Board of Canada Secretariat's Directive on Automated Decision-Making. Aimed at Canadian civil servants, it was created for the purpose of managing public sector AI delivery and procurement standards. So far, it has been used to complete four AIAs in Canada.

The AIA model consists of an online questionnaire divided into eight sections containing 60 questions on "technical attributes of the AI system, the data underpinning it and how the system designates decision-making". Impacts are ranked on a sliding scale from 'little to no impact', to 'very high impact' across a range of concerns, from individual rights and health and wellbeing to economic interests and the surrounding ecosystem. Once completed, each assessment is exported and uploaded to the Open Canada website.

The report cautioned that AIAs were not intended to replace existing regulatory frameworks but to complement those already used in the UK.

"This [process] is… proposed as one component in a broader accountability toolkit, which is intended to provide a standardised, reflexive framework for assessing impacts of AI systems on people and society," it said.

Through the pilot programme, the NHS is expected to support researchers and developers with information obtained through engagement with patients and healthcare professionals.

“Building trust in the use of AI technologies for screening and diagnosis is fundamental if the NHS is to realise the benefits of AI,” said Brhmie Balaram, head of AI research & ethics at the NHS AI Lab. “Through this pilot, we hope to demonstrate the value of supporting developers to meaningfully engage with patients and healthcare professionals much earlier in the process of bringing an AI system to market.

“The algorithmic impact assessment will prompt developers to explore and address the legal, social and ethical implications of their proposed AI systems as a condition of accessing NHS data. We anticipate that this will lead to improvements in AI systems and assure patients that their data is being used responsibly and for the public good.”

The report noted that it could not offer details on the Canadian government’s experience of the AIA process, nor any information of changes made as a result.

“Policymakers may be disappointed to find that AIAs are not an ‘oven-ready’ approach, and that this AIA will need amendments before being directly transferable to other domains. We argue there is real value to be had in beginning to test AIA approaches within, and across different domains,” it said.

“Policymakers should pay attention to how this proposed AIA fits in the existing landscape, and to the findings related to process development that show some challenges, learnings and uncertainties when adopting AIAs.”

The new pilot complements ongoing work from NHS AI Lab’s ethics team on ensuring datasets for training and testing AI systems are diverse and inclusive. It introduced the AI Ethics Initiative to support research and practical trials alongside wider efforts to develop and regulate AI-driven technologies in health and care, with the overall goal of preventing and addressing health inequalities.

Health Sites Let Ads Track Visitors Without Telling Them

Lily Hay Newman
February 6, 2022

All too often, digital ads wind up improperly targeting the most vulnerable people online, including abuse victims and kids. Add to that list the customers of several digital-medicine and genetic-testing companies, whose sites used ad-tracking tools that could have exposed information about people's health status.

In a recent study from researchers at Duke University and the patient privacy-focused group the Light Collective, 10 patient advocates who are active in the hereditary cancer community and cancer support groups on Facebook—including three who are Facebook group admins—downloaded and analyzed their data from the platform's "Off Facebook Activity" feature in September and October. The tool shows what information third parties are sharing with Facebook, and its parent company Meta, about your activity on other apps and websites. Along with the retail and media sites that typically show up in these reports, the researchers found that several genetic-testing and digital-medicine companies had shared customer information with the social media giant for ad targeting.

Further analysis of those websites—using tracker identification tools like the Electronic Frontier Foundation's Privacy Badger and The Markup's Blacklight—revealed which ad tech modules the companies had embedded on their sites. The researchers then checked the companies' privacy policies to see whether they permitted and disclosed this type of cross-site tracking and the flow of data to Facebook that can result. In three of the five cases, the companies' policies did not have clear language about third-party tools that might be used to retarget or reidentify users across the web for marketing.

"My reaction was shock at realizing the big missing pieces in these policies," says Andrea Downing, a coauthor of the study, independent security researcher, and president of the Light Collective. "And when we talked to some of these companies it really seemed like they just didn't fully understand the ad tech they were using. So this needs to be an awakening."

Downing and study coauthor Eric Perakslis, chief science and digital officer at Duke University's Clinical Research Institute, emphasize that while targeted advertising is a broadly opaque ecosystem, the tracking can have particular implications for patient populations. In the process of reidentifying users across multiple sites, for example, a third-party tracking tool could gather together information about a user's health status while also building a broader profile of their interests, profession, device fingerprints, and geographic region. And the interconnectedness of the ad ecosystem means that this composite picture can potentially pull in information from all sorts of web browsing, including activity on sites like Facebook. One classic example is the invasive targeted ads pregnant people and others consistently face based on marketer assumptions about their health status.

"The question in this experiment was 'Can patients believe the terms and conditions they agree to on health-related sites? And if they can't, do the companies even know that they can't?'" Perakslis says. "And many of the companies we looked at aren't HIPAA-covered entities, so this health-related data exists in an almost wholly unregulated space. Research has consistently shown that the flow of such information for advertising can disproportionately harm vulnerable populations."

The vast majority of users, of course, click through terms of service and privacy policies without actually reading them. But the researchers say that this is all the more reason to shed light on how digital ad targeting, lead generation, and cross-site tracking can erode user privacy.

“It’s entirely expected from my perspective that findings like this keep coming up for the category that I call 'health-ish' data that does not cleanly fall under the limited privacy protections that currently exist in US laws,” says Andrea Matwyshyn, a professor and researcher at Penn State Law and a former FTC advisor. “The evolution of terms of use when combined with privacy policies has created a murky picture for users, and when you try to analyze the data flows, you end up in this often endless spiral.”

The United States Federal Trade Commission established a Health Breach Notification Rule in 2009 that applies to health-related organizations not covered by the Health Insurance Portability and Accountability Act but has never taken an enforcement action under it. The agency gives examples of situations that could trigger enforcement, though, including one where a digital medicine company shares users’ medical information and mobile identifiers with an ad network without user consent.

The researchers focused on five consumer health-related companies: Color Genomics, Myriad Genetics, Health Union, Invitae, and Ciitizen. Invitae acquired Ciitizen in September, and the researchers found that the two companies went the farthest in their privacy policies to detail how they might use tracking technologies, including cookies and web beacons, that feed data to third-party services. Downing notes that both Invitae and Ciitizen could have gone into more detail about some of the specifics of their schemes, but overall they were clear that users could be subject to ad tracking on their sites.

Nonetheless, Invitae and Ciitizen are taking additional action as a result of the researchers' findings. “We are now in the process of suspending all of our ads on Facebook and removing trackers related to Facebook from our website in order to give the team time to fully understand, confirm, and eliminate any uses of data that could conflict with Invitae and Ciitizen policies or commitments,” Ciitizen's data stewardship and data-sharing lead, Deven McGraw, told WIRED in a statement.

Lauren Lawhon, Health Union’s president and chief operating officer, told WIRED in a statement that the company did not receive the researchers' disclosure about potential privacy issues until after their paper was published. She says that the company coincidentally conducted a major overhaul of its privacy policies throughout 2021, culminating in significant updates in December. When someone visits Health Union for the first time, they now see a pop-up to accept or reject data collection cookies and other tracking. Lawhon also notes that users can opt in or out of data sharing at any time, and the bottom of every Health Union community page now includes a “DO NOT SELL MY INFORMATION” link to surface these controls. Lawhon added that these changes came alongside "some improvements to how privacy management occurs.”

Myriad Genetics did not detail specific review or changes to its policies as a result of the findings, but it said that “no personal health information” from its quiz products is used to target individuals and that it complies with Facebook's health care advertising policies. Color Genomics says that it hasn't actively used two of the cross-site trackers (Leadfeeder and Nanigans) the researchers detected on its site in almost a year, and that it's continuing to look through the research findings.

For its part, Meta forbids organizations that use its activity trackers and other advertising and marketing tools from sharing health data with the social network. “We don’t want websites or apps sending us sensitive information about people,” the company writes on a resource page about sensitive health data. The company says it deploys automated tools designed to filter out any such data before it gets applied to serving ads.

Still, Meta's business model hinges on personalized advertising. In November, the company announced a “difficult decision” to remove thousands of sensitive ad-targeting categories related to topics like political beliefs, sexual orientation, religion, and race. The move also included removal of health-related categories like “Lung cancer awareness" and "Chemotherapy."

In its announcement about removing sensitive "Detailed Targeting" categories, Meta articulated the problem Downing and Perakslis examined in their research.

“The interest targeting options we are removing are not based on people’s physical characteristics or personal attributes, but instead on things like people’s interactions with content on our platform,” Meta explained. “We’ve heard concerns from experts that targeting options like these could be used in ways that lead to negative experiences for people in underrepresented groups.”

Downing and Perakslis consulted with the CERT Coordination Center at Carnegie Mellon University about the process for disclosing their findings. The group works with researchers to catalog software vulnerabilities and coordinate their public disclosure. But Art Manion, a CERT vulnerability analysis technical manager, points out that CERT and other organizations are only really set up to coordinate disclosure of specific software vulnerabilities and typically don't have structures in place to assess pervasive data leakages and notify relevant organizations about structural privacy issues. Instead, privacy researchers are often left to attempt an ad hoc disclosure process and then rely on companies across the digital ad ecosystem to have good intentions and make real changes.

“Right now there are almost no limits in place for what kinds of data companies can use to target their advertising, so that incentivizes them to collect as much as they possibly can," says Evan Greer, deputy director of the digital rights group Fight for the Future. “But the more data companies collect and store, the more likely—or rather, inevitable—it is that some of that data will leak in some way. Entities like the FTC should wildly scale up enforcement related to surveillance-based advertising.”

And while systemic privacy problems have persisted across the targeted ad industry, Downing and Perakslis emphasize that when it comes to vulnerable communities like patients and community organizers, there's a pressing need for clear policies and controls.