Algorithmic Self-Diagnosis from Targeted Ads

SARITA SCHOENEBECK, University of Michigan, USA
CAMI GORAY, University of Michigan, USA
YASEMIN GUNAL, University of Michigan, USA

People go online for information and support about sensitive topics like depression, infertility, death, or divorce. However, what happens when such topics are algorithmically recommended to them even if they are not looking for it? This article examines people's self-diagnostic behaviors based on algorithmically-recommended content, for example, wondering if they might have depression because an algorithm pushed that topic into their view. Specifically, it examines what happens when the sensitive content is not generated by users, but by companies in the form of targeted advertisements. This paper explores these questions in three parts. The first part reviews literature on self-diagnosis and targeted advertising. The second part presents a mixed-methods study of how targeted ads can enable self-diagnostic reactions. The third part reflects on the mechanisms that influence self-diagnosis and examines potential regulatory implications.

CCS Concepts: • **Human-centered computing → Collaborative and social computing theory, concepts and paradigms**.

Additional Key Words and Phrases: self-diagnosis, algorithms, sensitive, targeted advertising, privacy, companies

ACM Reference Format:
Sarita Schoenebeck, Cami Goray, and Yasemin Gunal. 2025. Algorithmic Self-Diagnosis from Targeted Ads. *Proc. ACM Hum.-Comput. Interact.* 9, 2, Article CSCW028 (April 2025), 19 pages. https://doi.org/10.1145/3710926

1 Part I: Introduction and Background

People have turned to the Internet to acquire instrumental and social support for decades. In 1997, Howard Rheingold wrote about the parenting group on The Whole Earth 'Lectronic Network, known as "The WELL", where parents talked about their children's illnesses and shared advice and support with each other [66].

With the widespread rise and adoption of online communities in the late 1990s and early 2000s, people turned to online communities for support on sites like Reddit and Facebook [30]. Today, people continue to rely on those sites as well as newer ones like TikTok (e.g. [10]). This reliance on online support is, of course, an extension of offline phenomena. Life would be difficult and lonely if we could not turn to each other, whether offline or online.

However, what happens when that advice is algorithmically delivered to users, whether they are looking for it or not? Media and scholars have raised some concern about self-diagnosis on TikTok, especially where people who are not seeking a diagnosis may be targeted for "algorithmic pseudo-diagnosis" [3, 13]. Gaeta [27] argues that the algorithms behind targeted ads serve to tell users what is "wrong" with them, and that this kind of "diagnostic advertisement" is a form of cure

Authors' Contact Information: Sarita Schoenebeck, University of Michigan, USA, yardi@umich.edu; Cami Goray, University of Michigan, USA, cgoray@umich.edu; Yasemin Gunal, University of Michigan, USA, gunaly@umich.edu.

ACM 2573-0142/2025/4-ARTCSCW028
https://doi.org/10.1145/3710926

and control. Users may see content recommended to them and engage in what Alper et al refer to as "platformed diagnosis", where a platform and a diagnosis are mutually constitutive [3].

While recognizing one's own symptoms and needs is important, overdiagnoses or misdiagnoses, as well as underdiagnoses, can lead to misleading recommendations and potentially harmful behaviors [60, 61]. Intervening in these cases is tricky; it is difficult to say with certainty which self-diagnoses are "legitimate" and we should take care not to be overly critical of self-diagnosis when the medical industry itself can be inaccurate and dismissive (e.g. [24, 31, 60]). Indeed, some research suggests that self-diagnoses can be equally valid as formal diagnoses [3, 22].

Even if there was some consensus that self-diagnoses from user-generated content online are harmful, it would be difficult to take strong action under current U.S. law. Section 230 of the Communications Decency Act protects companies from liability for most of the content that users post on their platforms [43]. In turn, companies' Terms of Service allow users to post a wide range of content. This is mostly good – we want people to enjoy freedom to express themselves. We also want to be very careful about restricting health information. What is health misinformation to one person may be life-saving advice to another and speech restrictions can be easily weaponized [53].

The point so far is that algorithmic self-diagnosis can be a concern, but we should not overcorrect by cutting off basic rights and societal goods. We now consider an alternative variant to this issue: what about when companies themselves are the ones creating and recommending content that encourages self-diagnostic behaviors? If a company persistently markets content to you recommending depression treatment, and you start to wonder if you have depression, is this okay? This hypothetical is worth examining – many thousands of other companies target advertisements to us in our feeds all day every day, many of which contain products, services, and commentary about sensitive topics [42].

This article examines *algorithmic self-diagnosis*, where people might internalize something about themselves based on content that is algorithmically recommended to them. It focuses specifically on content that is targeted to people via advertisements, which has distinct characteristics compared to user-generated content in news feeds. It also has distinct characteristics compared to other online behaviors like search or social media. Rather than identifying symptoms and looking for explanations, self-diagnosis from targeted ads treats the algorithmic recommendation like a symptom, inducing people to look for explanations.

This article proceeds in three parts. The remainder of Part I reviews prior literature in Internet self-diagnosis and in targeted advertising. Prior work has examined mechanisms that explain self-diagnostic behaviors in Internet searches and online health communities. However, self-diagnostic behaviors have not been theoretically or empirically linked to targeted ads. Part II presents an empirical study of people's reactions to target ads. We draw from mixed methods analysis of a survey and digital ethnography of reactions to targeted ads. The analysis focuses on people's self-diagnosis reactions and reflections, especially for ad topics that are sensitive in nature. Part III examines the mechanisms that can influence self-diagnosis. It then discusses potential regulatory protections and the tradeoffs between user autonomy – often a priority in the field of Human-Computer Interaction–and regulatory protections. The article makes the following contributions:

- Shows what kinds of targeted ad topics are likely to induce self-diagnostic behaviors (Part II)
- Reveals what kinds of self-diagnostic reactions people have to targeted ads (Part II)
- Explores the mechanisms that shape self-diagnostic reactions to targeted ads (Part III)
- Examines regulatory considerations for self-diagnosis and targeted ads (Part III)

1.1 The Internet and Self-Diagnosis

Online health communities (such as Facebook groups or cancer support websites) are important sources of informational and social support [29, 51, 52]. They can connect people with uncommon illnesses across geographic distances and can help people identify what their symptoms might be indicating. They can also be anonymous, allowing people to discuss sensitive topics that they may not want to discuss with people that they know [5]. When not anonymous, they can allow users to build trust and rapport with each other, which may be especially helpful when navigating long-term illnesses [4]. Despite these benefits, online health communities also introduce many concerns around misinformation and misdiagnosis [76]. Information is rarely vetted and can be inaccurate or misleading. Recommendations can also induce risky behaviors and outcomes, such as eating disorder communities that praise harmful behaviors [73].

Self-diagnosis refers to the process of identifying symptoms and conditions in oneself without the expertise or input of a professional. As Tik-Tok has risen to prominence, there has been growing awareness of people – especially young people – engaging in self-diagnosis of mental health diseases based on TikTok content they see. Content creators post videos where they have self-diagnosed their mental health and list the characteristics that led them to this self-diagnosis. Some medical professionals also act as influencers and content creators, which can create an aura of legitimacy without the oversight of the traditional medical industry [36]. Most sites have little vetting of user authority or expertise: a trained expert is typically treated the same as an everyday person online. Influencers may participate in communities where people offer and share support for their diagnoses, such as #depressiontok on TikTok, a portmanteau for depression and TikTok. As a result, influencers can have out-sized influence, regardless of whether they have the credibility and legitimacy to do so. The sites that do establish vetting practices, such as Wikipedia, have largely limited the spread of misinformed and unvetted content or users [8, 46], but most digital platforms prioritize user speech rather than vetting content.

Self-diagnosis can be a powerful act – the ability to find a label or an explanation for a struggle can be validating and supportive [3, 22]. This is especially important given that many people do not have access to healthcare services to attain formal diagnoses, or who may be overlooked by formal healthcare providers. Alper et al. and Eagle et al. present stories of self-diagnosis from autistic people and people with ADHD, respectively, highlighting how self-diagnosis can be a precursor to a formal diagnosis, or it may be valid and legitimate independent of any formal diagnosis [3, 22]. They also describe how attaining diagnoses can be difficult, whether due to being ignored or dismissed by medical providers or not having access to them.

Some of the benefits and risks of online health communities for identifying health conditions are also observed in Internet search. Myriad papers have been written about "Dr. Google", describing patients' practice of looking up their symptoms online and then bringing them to the doctor to suggest what they may have [45, 49, 74]. It is worth noting that the Internet did not create this phenomenon: people have long expressed their beliefs about what their symptoms may be indicating. Early doctor's notes from 100 years ago wrestle with balancing patient self-reports and declarations with the doctor's own expertise and recommendations [40]. Like with online health communities, the practice of searching for information to make sense of one's own experiences and symptoms is powerful and empowering. People whose illnesses may have gone ignored, dismissed, or misdiagnosed can find new sources for information and support.

However, Internet search introduces many similar concerns as other digital health behaviors. In 2009, White and Horvitz embraced the term "cyberchondria" to refer to Internet users' escalation of concerns based on their symptoms and web searching behavior [84]. Their studies found that escalation of concerns is associated with how much medical content people see, how much escalating language is there, and a person's predisposition towards escalation versus looking for other possible explanations [84]. When searching for health information, people are more likely to focus on alarming content (e.g. a heart attack) regardless of its location in search results or position on a page [85]. They also find evidence that users engage in confirmation bias and anchoring-and-adjustment behaviors [86]. Specifically, people who search for health information tend to hover on pages that align with their beliefs and to spend less time on those that do not. They also find that people who are confident in their beliefs are less likely to revise those beliefs despite what search results may produce.

An alternative to relying on search results is to engage with more personalized health data and recommendations. Tools like Fitbit, Apple Watch, or Strava are all designed to collect data on personal behavior and display it back to users. This genre of app, which supports data collection and personal reflection, has often been referred to in the scholarly literature as "personal informatics" or "quantified self" [23]. Though empowering people to make choices about their health has typically been lauded, tracking apps can be cognitively burdensome to use and make sense of (see [17, 75]. Along with the rise of personal informatics, the 2010s saw a rise in interactive, diagnostic apps, sometimes described as "direct-to-consumer". These apps tend to use algorithms to derive recommendations from user data, and sometimes use computer vision for photos (e.g. dermatology). Some also use online symptom checkers or crowdsourcing. They allow fast and convenient personalization but are also subject to diagnostic error. They are not subjected to FDA review because they are seen as an important innovation that should not be regulated [57].

1.2 Targeted Ads

Advertisements are about selling products, not helping people [2]. Targeted advertising is a form of advertising where ads are personalized, or targeted, based on a person's identity and behaviors. For example, a person who spends a lot of their time on ESPN.com and NFL.com may find that their Instagram profile is targeted ads for team jerseys or for brands of beer. This is in contrast to contextual advertising which presents advertisements that are contextually appropriate. In that case, a person who visits NFL.com may see jersey or beer ads on that site while browsing but would not necessarily see them on other sites or apps they visit. Targeted ads require sophisticated third-party tracking across sites and data brokers who transact in consumer data. Targeted ads emerged over the past 15 years as advances in machine learning facilitated large scale data collection and analysis, allowing advertisers to target consumers at a much more granular level than in previous generations (e.g. television ads). Demographic targeting still exists but instead of by television channel and time of day (e.g. to stay at home mothers), advertisers can direct ads to relevant segments by digitally targeting audiences, keywords, or content (e.g. Spotify's category of users who attend university).

Targeted advertising works through the use of what are known as "cookies", small bits of code that are embedded in websites and can request or store information on a user's browser when they visit the site. First-party cookies are stored directly on a website and track users when they visit that website. First-party cookies can be useful – we do not have to login to every site every time we visit it because our login information can be stored in a cookie. Third-party cookies are created by another entity, not the host site, and are embedded in code on a website if there is an arrangement between the website and the third party broker to do so.

An article from 2012 by Ur et al. lays out the landscape of targeted ads concerns [78]. They outlined the privacy tradeoffs (targeting was useful but also creepy), the problems with consent (people did not really know how to avoid targeting), and a limited understanding of data practices (people did not understand how targeting is done) [78]. More recent bodies of work have been two-fold: one line of work has looked to develop themes and patterns to explain the psychological underpinnings, user explanations, and potential harms of online behavioral advertising (e.g. [12, 34, 65, 81, 87, 88]).

A second line has focused on specific groups associated with targeting, including young people's reactions to targeting [33], queer people's experiences with targeting [67], and people with histories of eating disorders' experiences with targeting [28]. A theme across this second body of literature is that targeted ads reify dominant ideas about who or what is "normal," whether it be adulthood, heternormativity, thinness, or the myriad other categories and characteristics that comprise us.

In response to these concerns, the European General Data Protection Regulation (GDPR) began requiring the websites to allow users to opt-out of third party cookies.[1] This has increased attention to the problem, but the burden of opting-out on every website is prohibitive and most people engage in illusory "consent" to data tracking because it's too much work to do otherwise [79]. At the same time, many sites have ramped up user controls to opt-out of categories of ads. Google allows users to toggle off alcohol and gambling ads specifically, while Facebook offers users to select "Show less ads about this topic" for a range of categories.[2] However, the concern is that, like notice and choice regimes, these opt-out features require users to know how to set their preferences and willing to set them across sites, and they do not any user feedback as to whether the settings are working.

Regulatory protections of privacy online have largely focused on the concept of personal data [18, 70]. Personal data refers to the myriad bits of information that are linked directly to us, such as our names, friends, likes, groups, purchases, and even our scrolling. Some data are deemed sensitive, at which point they are subject to heightened protections [63, 71]. For example, financial data is regulated by laws such as the Gramm-Leach-Bliley Act and requires more advanced security (e.g. for online banking).[3] Similar principles apply to health data, personal data like social security numbers, educational data, and biometric data. However, scholars have raised concerns that the concept of sensitive data is arbitrary and incoherent [63, 71, 72]. They argue it is not clear why some data enjoys heightened protections and others do not, and even if it was clear, providing protections for data does not provide protections against sensitive inferences. Inferences from people's digital data can be unpredictable or troubling and there is little oversight over those predictions [83]. Such inferences may also violate people's right to intimacy, dignity, and self-determination [18].

2 PART II: Empirical Study

We conducted a mixed methods study pairing an online survey with observational data. The online survey asked participants about their comfort with targeted ads and why they were comfortable or not. The observational study involved analysis of posts and discussions about targeted ads. The mixed methods approach allowed us to triangulate survey data with observational data. We used the observational study to inform the survey design, especially to determine what kinds of ads to ask about in the survey. We then continued with observation to make sense of and interpret survey results.

2.1 Online Survey

We conducted an online surveywhere we asked participants about their comfort with various ad topic types. This was part of a broader series of surveys on targeted ads and sensitive topics. A prior study focused on people's comfort with sensitive ad topics [68]; the current project focuses on free response questions about topics participants were not comfortable with and why. Participants first rated their comfort with being targeted ads for different topics (e.g. fertility, divorce) and in different contexts (e.g. which company was advertising, where the ad was seen). We selected the topics to focus on based on prior work and our observations from discussions about targeted

[1] https://www.consilium.europa.eu/en/policies/data-protection/data-protection-regulation
[2] https://support.google.com/My-Ad-Center-Help/answer/12155260?hl=en, https://www.facebook.com/help/353660662271696
[3] https://www.ftc.gov/business-guidance/privacy-security/gramm-leach-bliley-act

Cancer treatment	Infertility treatment	Bankruptcy lawyer
Weight loss products	Burial or cremation services	Baby products
Hair loss treatment	Birth control	Divorce lawyer
Depression treatment	Alcoholic beverages	Dating websites
Anxiety treatment	Tobacco or e-cigarettes	Political candidates
Eating disorder treatment	Gambling websites	Paper towels
Therapy or counseling	Weed or marijuana	Bicycles
Your favorite foods	Your favorite type of shoes	Headphones
Religious services	Sexual enhancement products	Can openers

Table 1. Ad topics participants were asked about their comfort with and why.

ads online [6, 28, 67]. We brainstormed a longer list of topics and then refined it after discussion and pilot testing. The topics ranged from physical and mental health to finances to relationships to behaviors. We chose these because they can be personal and sensitive topics that people may not be comfortable disclosing with others. We also chose some topics that we expected not to be sensitive. This allowed us to corroborate whether people were not comfortable with all ad topics, even non-sensitive ones, or was their discomfort specific to some types of ads. The control topics included favorite foods and favorite types of shoes , headphones, paper towels, bicycles, and can openers. We selected items that should be relatively accepted and that everyone should be familiar with. The 28 topics are not comprehensive which is unrealistic; they were designed to draw out prominent patterns that may be especially important for scholars, industry, and regulators to consider. Our pilot testing suggested that having participants rate all 28 topics would lead to fatigue, so the survey presented a subset of 10 of the 28 different ad topics to participants selected randomly and presented in randomized order (see Table 1).

Participants were given the prompt: *Websites track your online behaviors, such as websites you have visited, Internet searches you have made, your purchase history, and your cell phone location. Websites then use that data to target ads to you to keep the websites and services you use free of charge. How comfortable are you with websites targeting the following ad topics to you?.* After rating all the topics, on a new page, for each ad topic that was rated "Somewhat uncomfortable" or "Extremely uncomfortable", participants were asked why it was uncomfortable for them. Where they rated the topic as uncomfortable, they were asked to explain why.

Participants were recruited on the online platform Prolific in early 2023 (see Table 2). Participants had to be 18 or older and in the U.S. and speak English. We sampled for diversity in age, gender, ethnicity, sexuality, religion, health, disability, and behaviors (e.g. alcohol use). All participants were shown a consent form before proceeding to the survey. After data cleaning, there were a total of 1,003 participants. Completion time was between 4-5 minutes for most participants and compensation was over $15/hour on average. This research was exempted from full review by the Institutional Review Board. This study presented low risk to participants – the topics we asked about were topics that could be seen in everyday Internet browsing or might come up in everyday social life.

We analyzed data using an inductive and thematic approach where two of the coauthors read and discussed the data and analysis [77]. To begin, we went through the all of the data in multiple passes and discussed potential topics and themes of interest. After developing a draft codebook based on our reading and discussions, we pilot tested it by having one coauthor code a random sample of 100 responses and then discussed points of uncertainty and potential disagreement with the second coauthor to establish reliability [55]. After finalizing the codebook, one coauthor coded the entire dataset with input from another coauthor. Self-diagnosis was a theme that we observed in the data and we decided to focus on it because it was important and not well documented in prior literature. We then revisited the coded data with a focus on data that was coded as self-diagnosisWe developed categories self-diagnostic behaviors and then report on themes from that analysis.

Age M=39; SD=14.85	White 71%	One or more disability 83%
Woman 46%	Black 14%	One or more children 43%
Man 45%	Hispanic 9%	In a relationship 68%
Non-binary 7%	Asian 9%	More politically left 58%
Transgender 7%	Native/Indigenous 5%	More politically right 23%
LGBQ 30%	Middle Eastern 3%	Religion is very or somewhat important 43%

Table 2. Participant demographics.

2.2 Digital observation

We conducted a yearlong digital observation of self-diagnosis related behavior and targeted ads online. Our approach borrowed techniques from digital ethnography in its length of engagement and reflexive practices, but did not follow people or communities persistently in ways that might be expected of an ethnography. In our observations, we focused on online platforms like Twitter, Reddit, Instagram, TikTok, and YouTube because we were familiar with those platforms as users and researchers. We also chose them because prior research has focused on health related posts on those platforms, giving us prior insights to build on (the body of work is extensive, see a few examples e.g. [6, 14–16, 20, 58, 62]).

As is typical with digital observation and qualitative analysis, our data collection relied on purposively selected content. We browsed Twitter, Reddit, Instagram, TikTok, and YouTube discussions about ads and sensitive topics. We searched for relevant content using Twitter and Reddit's search interfaces, using keyterms "ads" or "advertisements" with either the condition, such as "depression", or a company, such as "BetterHelp". This search strategy was not intended to collect a comprehensive or representative sample of discussions. It allowed us to focus on posts that were likely to be relevant for our inquiry. Our focus was on qualitative interpretation and we did not attempt to quantify the presence of self-diagnostic related ad content in this study. Our data collection involved note-taking of phenomena and collecting screenshots of advertisements in some cases.

Our analysis was conducted throughout the yearlong process. In the early stages, we used our observations in tandem with literature reviews to develop the ad topics used in the survey. In the later stages, we use the observations to support qualitative responses from the survey. As new themes emerged inductively, we also iteratively revisited those results to accumulate and compare findings across data sources [56].

We present our results grouped into prominent themes. We use survey quotes and online content to illustrate themes. The online content is edited to preserve anonymity because online users did not expect or agree to their data being used in research and because the topics are relatively sensitive [25]. To align with user preferences, we also did not look at profile information and tried to ensure that content could not be linked back to the user [25]. The goal is not to identify or call attention to particular users, but to shed light on an online phenomena. As emphasized by Fiesler, et al. [25], the ethical consideration we prioritize in this work is not whether the data is public (it is), but how to minimize risk or harm to participants and ensure that benefits of the research outweigh those risks or harms.

2.3 Results

Results are structured around three themes. The first theme focuses on what kinds of ads can trigger self-diagnosis. The second focuses on reactions to receiving the ads. The third focuses on the role of algorithms in fueling self-diagnosis reactions.

2.3.1 Ad topics that can trigger self-diagnosis. There were four overarching themes in the types of ad topics that led to self-diagnostic reactions: physical health, mental health, reproductive health, and relationships (see Figure 1). Physical health included cancer treatment and hair loss treatment. For both of these, participants reported that receiving an ad about it made them wonder whether it was something they should be thinking about. One participant said of cancer treatment ads, "I would worry that my previous search history or behavior indicated I was a target for the ad or needed the product."

Mental health included depression treatment, anxiety treatment, eating disorder treatment, and therapy or counseling. Many participants reported that they received a lot of mental health ads and that those ads triggered self-diagnosis. Companies like BetterHelp and Headspace advertise pervasively and people expressed frustration with the frequency of these ads, but also insecurity and anxiety about why they might be getting them. Some ad topics spanned both mental and physical health domains, including burial or cremation services and weight loss topics. Weight loss product advertisements tended not to induce self-diagnosis about weight or weight loss specifically – most people may be already aware of their body and how they feel about it – but they fueled internalization of weight related issues. Weight loss ads were prevalent among participants. The number and range of weight related ads some people receive may make it difficult *not* to internalize them.

Infertility treatment ads spanned physical, mental, and reproductive health categories in our analysis. This was primarily about birth control and baby products, topics that may be magnified in people's awareness due to some prominent news stories linking search behavior to pregnancy diagnoses (e.g. [38]). It may also be that some people harbor persistent fears about reproductive health. This might have been especially salient during the time of the survey in 2023 soon after Roe v. Wade was overturned, which could amplify the effects of targeted ads about that topic.[4]

Addiction-related ad topics did not elicit self-diagnostic responses in our data. It may be that people who are heavy or problematic consumers of alcohol, drugs, or gambling already know that about themselves. It may also be that self-diagnosis is more relevant to health and relationships than to behaviors, even if those behaviors have underlying mental and psychological characteristics.

2.3.2 Reactions to receiving the ads. Participants described how seeing ads for some topics induced anxiety for them. This was especially the case for health related ads, both mental and physical. Many people talked about how eating disorder and weight loss ads could make them feel anxiety about what they were or were not currently doing. One said, for example, "I also don't deal [well] with eating disorders and showing me ads relating to it might make me double think it and just make me feel anxious." Another said of eating disorder ads, "Although helpful, it would increase my anxiety to see the help advertised."

Participants were also concerned about receiving ads for cancer which they worried could signal something was wrong that they did not know about: "I don't have cancer or deal with it... hinting that I might with these ads might make me anxious though." Another said, "I get anxious about my health, so commercials about cancer treatments don't help my anxiety."

Participants expressed anxiety about the idea of receiving ads for burial or cremation services, which made them feel like this was something they *should* be thinking about: "scared they think

[4]https://www.supremecourt.gov/opinions/21pdf/19-1392_6j37.pdf

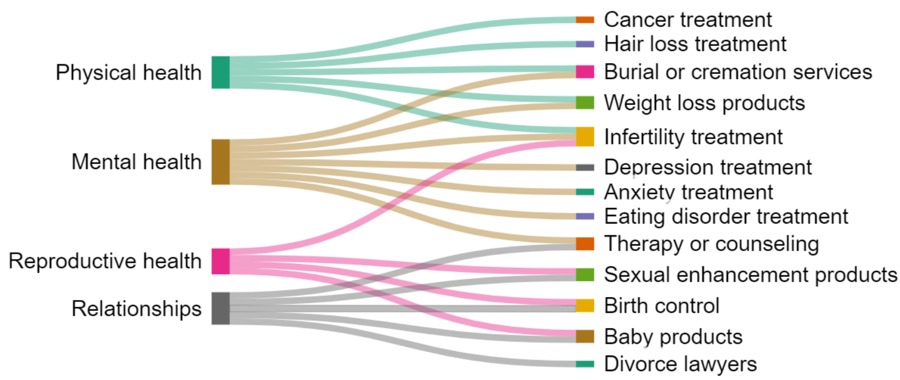

Fig. 1. Ad categories and topics with self-diagnosis reactions.

im old enough to have this on my mind." Participants also mentioned anxiety related to infertility ads, e.g. "thanks I'm scared now." They also expressed anxiety that death might be coming sooner than they imagined: "I don't have a need for these services and it would make me feel like the advertisers think i will be dying soon."

The concept of internalization describes the process of an idea moving from the outside to inside of the mind. To internalize something is to incorporate into oneself, identity, beliefs, or values. Some participants internalized the topics associated with targeted ads, causing them to think something different about themselves that they otherwise would not have. These spanned many ad topics. For example, about weight loss, one said "Having weight loss products targeting at me would make me feel like I need to lose weight." Another said that ads about Burial or cremation services "makes me feel I am getting older." On the same topic, another said "About to have an existential crisis knowing that my sponsored ads are aging me 20 years."

One topic some of them mentioned insecurity about was hair loss. For example, one said "Hair loss is a sign of advancement in years, so this is not something that I need to be reminded of." Another asked, "That being said, why are my ads all hair transplant-related and hair growth products??"

The frequency of ads triggered negative reactions from participants, especially related to mental health treatment ads. One participant said, "I think anxiety ads make me feel like it's very prevalent and maybe I suffer from that too." One online user said, "Do all of you get @betterhelp ads pretty much all the time or is the algorithm saying its worried about me?" Another said, "If I see another Betterhelp ad on YouTube i'm gonna actually need therapy. Nothing against Betterhelp here good god it shows on EVERY SINGLE VIDEO I CLICK. FOR THE PAST MONTH." Participants in our survey also mentioned BetterHelp specifically, e.g. "I'm sick of these @betterhelp ads on my feeds trying to gaslight me into thinking I need therapy."

2.3.3 How algorithms can fuel self-diagnostic reactions. Many of the self-diagnosis reactions participants had were related to the algorithmic nature of ad delivery. This included the personalization associated with algorithms and the relational nature of algorithms.

Participants often reacted negatively to the personalized nature of ads. They believed there must be some reason for the ad targeting. For example, "I don't have any major anxiety issues, but if my internet behavior is making some Instagram algorithm conclude that I need anxiety treatment, that would probably make me more anxious." For mental health, many felt it was predatory to target

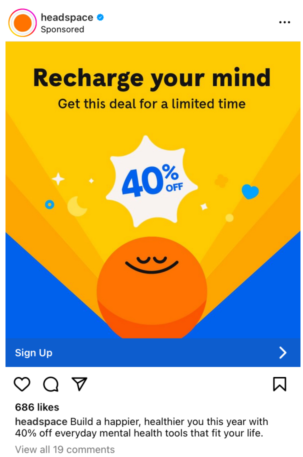

(a) Mental health

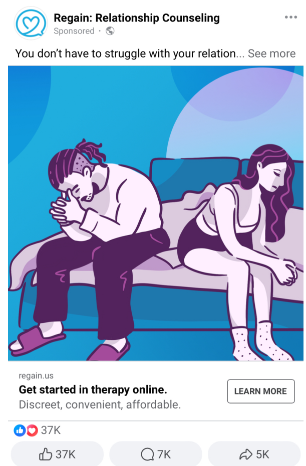

(b) Relationship

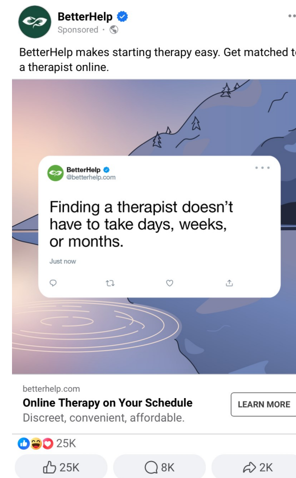

(c) Mental health

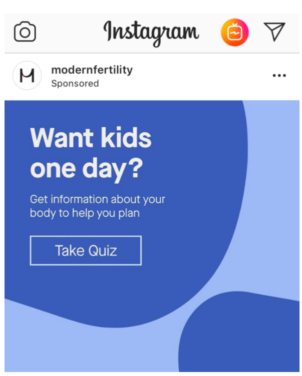

(d) Fertility

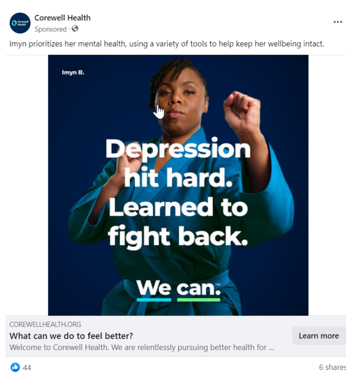

(e) Depression

(f) Hair loss

Fig. 2. Example targeted ad topics.

ads. "If a company is offering therapy services they are aware that the people they are targeting with adverts are struggling at the moment, whether it be anxiety, depression, OCD, trauma etc. Their advertisements induce panic attacks in some people, including me."

Algorithmically delivered ads induced insecurity in some participants. For example, one said about weight loss ads, "I would feel insecure and wonder why this was targeted towards me." Another said about sexual enhancement product ads, "Having this targeting me would make me feel like I'm lacking." Some participants signaled insecurity related to aging and death, saying things like "Do they know something I don't??" in relation to ads for burial or cremation services. Another simply said, "They think I'm about to die." One said about depression ads, "Seeing these ads would make me ponder [whether] or not I have issues with depression." Another said about anxiety treatment ads, "I don't have anxiety but these types of ads are so general that they make you start to think you do have anxiety." One participant expressed discomfort with ads that they felt were trying to diagnose them, "I don't feel comfortable seeing ads that feel like they're trying to diagnose me."

One of the more common topics where participants internalized was related to infertility. Many participants worried that the algorithm might know something they didn't about their fertility. One noted, "this would make me worry that something I have been eating or doing is effecting my fertility without me knowing." Another said, "If I saw ads like this, it would make me think something was wrong with me physically." Some of them turned to sarcasm, especially in the context of fertility: "Since turning 30, I've noticed a significant increase in the frequency of targeted ads about infertility treatments. This isn't concerning at all.." Multiple people talked about ads they get on different platforms and what it might mean: "I always get infertility ads on my instagram and it makes me think"; "I'd worry that google somehow picked up infertility before I know"; and "What does YouTube know that I don't?!" They also expressed insecurity if they went from receiving no ads about infertility at all to receiving them frequently, believing there should be some explanation for a sudden change.

Participants talked about the relational nature of algorithms, particularly in the context of divorce ads. They felt that divorce ads would mean their partner was looking up divorce, or that their partner would think *they* were looking up divorce ads. One said "It would potentially make me think my significant other was looking at divorce lawyers." A person online noted, "Deeply concerned about how every ad Ive gotten the last two hours has been about divorce." One said that since they were not considering divorce, "it makes me think that my husband might have been googling for a divorce lawyer on my device." Another said they would not want their partner to think they were going to divorce them by seeing an divorce ad. They also noted the relational nature of ads based on what others might think about them. "I don't want the world to know that I need to drop five pounds. Sending me weight loss product ads makes me feel that others think I am fat." Though people usually don't see other people's ads, some participants felt others might know about them anyway. In a few cases, they explicitly mentioned that people might see ads over their shoulder at home or at work.

3 Part III: Discussion

Our results reveal the types of ads that can trigger self-diagnostic reactions and the nature of those reactions. This section has two subsections. The first explores the mechanisms that might influence self-diagnostic reactions to targeted ads. The second section explores what might be done about algorithmic ad self-diagnosis. It considers what regulatory protections might be considered and their strengths and limitations.

3.1 The Mechanisms that Influence Algorithmic Self-Diagnosis

3.1.1 Prevalence Inflation. Participants engaged in self-diagnosis in reaction to the frequency of content showing up in their feed. Algorithms can inflate the amount and kind of content that is recommended to a user, often reinforced by user behaviors that suggest interest in that content. This can create the perception of prevalence of a phenomenon. When we see content a lot, we think there must be a lot of that content (and that other people are also seeing a lot of it). The term prevalence inflation was used in 2012 in the context of mental health and epidemiology to depict how mental health disorder counts might be inflated due to differences in counting [64]. Specifically, inclusion of a larger number of disorders in a count can lead to higher overall prevalence, even when the prevalence of no specific disorder has increased [64].

Prevalence inflation was used more recently in a 2023 article that hypothesized that greater awareness of mental health can have dual effects [26]. On one hand, mental health awareness efforts lead to more accurate reporting of previously under-recognized symptoms which is a beneficial outcome. However, awareness efforts might also lead some individuals "to interpret and report milder forms of distress as mental health problems" [26]. Alper et al. discuss the challenges in navigating self-diagnoses in online environments, which can be both validating of a self-diagnosis but can also lead to accusations of illegitimacy of a self-diagnosis [3]. This article does not focus on addressing the challenges around legitimacy of self-diagnoses. However, it highlights the role that targeted advertisements can play in fueling those challenges, with little visibility or guardrails into how it may be impacting people and communities exposed to that content.

Prevalence inflation also relates to the availability heuristic, a principle often turned to in HCI, which refers to the ease with which one can bring to mind an event [69]. This is a cognitive bias that allows people to make fast assessments. It works by relying on mental shortcuts, such as events that happened recently or that were particularly memorable. The availability heuristic can be a useful cognitive tool – it allows us to see and remember things that may be useful for future decision-making [54]. It also allows us to accomplish mundane tasks like selecting an option from a dropdown menu. However, it can be used by advertisers to encourage people to purchase products and services. While this is accepted practice within advertising (e.g. billboards on highways), it is more ethically questionable to show people ads that may induce harm because of the targeted nature.

3.1.2 Information Asymmetry. Participants expressed discomfort with targeted ads because they did not know why they were being targeted. Ads are targeted through a complex set of data sources, including profile information, demographic information, behavioral data, location data, and scrolling data, among other types. Targeting is also relational. Targeting assumes that people will be interested in things that other people like them are interested in and uses group and population level behaviors to infer those interests [82].

Information asymmetry refers to when one party has more information than another, which the first party can typically use to their advantage. Information asymmetry is used to explain markets where sellers have more expertise than buyers, such as car dealers, health care, or banking [21, 50]. Most consumers have little expertise into how these professions or industries operate which makes it difficult for them to engage in transactions with complete or sufficient information. Information asymmetry also exists when technology is involved – between platforms and users, between AI and users, and between advertisers and consumers. In the case of targeted ads, users have little idea about why they may be targeted an ad because targeted advertising is largely a black box.

Like with many AI systems, auditing targeted advertising models on a large scale to understand what data goes into them and how ads are delivered is nearly impossible. It is technically possible to audit ads [47], but doing so on a large scale requires company access which they do not provide. The concern with people not understanding why ads are targeted at them is that they may behave in ways that are not in their best interest, such as believing they have a condition because a treatment is advertised to them. Jack Bandy refers to taking advantage of users with a health history to advertise more of the same as *algorithmic exploitation* [9]. This could be combated by providing more information to users about why they are seeing an ad. There could also be standards for what can be advertised and how often.

3.1.3 Perceived Legitimacy. Our results suggest that participants perceive some amount of legitimacy in algorithms. They referred to "the algorithms" as knowing something about them (that perhaps they didn't even know themselves). Perceived legitimacy refers to the belief that actions of an authority are appropriate and trustworthy. Legitimacy can be conferred on any authoritative agent – a person, government, court, or business. When people perceive an entity to be legitimate, they are more likely to trust and obey it. Legitimacy has been an important principle in criminal legal systems. People are more likely to obey laws and accept judicial decisions if they believe them to be legitimate [80]. And the opposite is true – perceived illegitimacy can lead to breakdown in order. Conferring authority to the algorithm requires "algorithmic legitimacy" – the belief that the output of the algorithm carries some kind of decision-making power.

People have historically expressed trust in machines, viewing them as objective and neutral, unlike humans who are biased and subjective [44]. Though we have much more nuanced understandings of sociotechnical systems now, there is still a lot of trust instilled in AI recommendations which are often trusted similarly to humans (e.g. [7]). From policing to traffic control to healthcare recommendations, AI predictions and recommendations drive a lot of decision-making. Ads may carry perceived legitimacy if people believe the algorithms that drive them and the industry that approves them are trustworthy.

One important consideration with the concept of perceived legitimacy is whether people *should* trust authorities. In important contexts like health, policing, and courts, there are biases and disparities that create unequal and discriminatory outcomes. There is also evidence of the same phenomenon with ads. In 2019, The Department of Housing and Urban Development sued Facebook for discrimination in housing advertisements because they allowed advertisers to target by race, religion, and national origin [1]. However, ads for mental health, physical health, reproductive health, and relationships are largely unregulated. Whether ads are delivered fairly, ethically, safely, or responsibility is an open question.

3.2 Exploring Regulatory Protections

This section explores potential protections for users and considers the benefits and risks of implementing such protections. It focuses on the targeted nature of ads and whether privacy protections are necessary. It also focuses on mental and physical health, which were major themes in our findings.

3.2.1 Privacy Protections. Participants in our study expressed concern about the targeted nature of ads. They expressed discomfort with algorithms "knowing" too much about them. Scholars have characterized this kind of algorithmic knowing through the lens of the "algorithmic crystal" where algorithms can shape how we view ourselves [48]. Algorithms are not only tools for self-knowledge and the co-production of knowledge, but users may imagine them as providential and spiritual in their ability to shape users' experiences [19, 41]. While the analytical lenses used to make sense of algorithms and self-concepts may vary, there is a theme across the literature that algorithms shape how we see ourselves [11, 39].

Though concerns have been raised for over a decade, in recent years, there has been growing regulatory proposals to address data collection and brokering practices, including the General Data Protection Act, the California Consumer Privacy Act, the Federal Trade Commission (FTC) Privacy and Data Protections, the White House Office of Science and Technology Policy, and the Consumer Protection Financial Bureau.[5] Some of these regulations specifically call out targeted advertising, such as the FTC's charge against Twitter.[6] Others make more general recommendations, which can be applied to regulate targeted ads. For example, the White House Bill of Rights says, "Designers, developers, and deployers of automated systems should seek your permission and respect your decisions regarding collection, use, access, transfer, and deletion of your data in appropriate ways and to the greatest extent possible; where not possible, alternative privacy by design safeguards

[5]https://gdpr-info.eu/; https://oag.ca.gov/privacy/ccpa; https://www.ftc.gov/news-events/news/press-releases/2022/05/ftc-charges-twitter-deceptively-using-account-security-data-sell-targeted-ads; https://www.whitehouse.gov/ostp/ai-bill-of-rights/; https://www.consumerfinance.gov/about-us/newsroom/remarks-of-cfpb-director-rohit-chopra-at-white-house-roundtable-on-protecting-americans-from-harmful-data-broker-practices/

[6]https://www.ftc.gov/news-events/news/press-releases/2022/05/ftc-charges-twitter-deceptively-using-account-security-data-sell-targeted-ads

should be used. Systems should not employ user experience and design decisions that obfuscate user choice or burden users with defaults that are privacy invasive."[7]

One outcome of this impending regulation related to targeted ads is an industry-wide shift away from third-party targeting. Instead of allowing third party cookies, Google and Amazon have announced they will aggregate and analyze user behaviors locally onsite and only communicate high level topics and interests to third parties.[8] This shift decreases privacy threats associated with third-party cookies, an industry which has been difficult to study and has gone largely unregulated. However, doing so consolidates data collection and processing within a few big companies that dominate the ad ecosystem, including Google, Amazon, Meta, and Apple (see [68]). If participants are uncomfortable with companies knowing too much about them, is it helpful for fewer companies to know about them, but those few companies to potentially know more about them? The decline of third party cookies is promising, but it does not mean we are not being tracked and targeted to anymore, and many of the regulatory protections being proposed may still apply to emerging advertising technologies.

3.2.2 Health Liability. Many of the ad topics that induced self-diagnosis in our study related to physical and mental health. Practicing health in the U.S. is heavily regulated, especially between medical providers and patients. There are myriad laws dictating what medical providers can or cannot do or say, and what rights patients have.

These regulations include the requirement for medical licensing to practice, medical liability law that protects against malpractice, informed consent laws that protect patient autonomy, and fiduciary responsibilities that require doctors to put their patients' interests ahead of commercial ones [36, 37]. A physician's "duty of care" stipulates that, if they accept a patient, they will attend, diagnose, refer, treat, and instruct the patient in a reasonable manner that would be consistent with other physicians.[9] Failing to do so can violate both ethical commitments (e.g. the Hippocratic Oath) and legal ones (e.g. malpractice).

Targeting advertising of health topics largely falls outside of that scrutiny. Companies can advertise a range of health products and services (e.g. the range of examples in Figure 2). However, if advertisements are for health care services that offer low quality or pseudoprofessional advice, there may be cause for regulatory protection. For example, the FTC's Health Products Compliance Guidance protects against false or misleading advertising claims about the benefits or safety of dietary supplements or other health-related products.[10] While people's interpretations or experiences of advertisements is not generally subject to oversight because then any ad could be unreasonably accused of causing harm, a question is whether the targeted nature of ads and the algorithmic amplification of such ads into a newsfeed changes that calculation. Seeing a billboard for mental health treatment on the freeway is acceptable, but is seeing repeated mental health ads that are algorithmically targeted to a person? Neither the law, nor ethical discussions, have discussed this much yet, but perhaps they should.

A related consideration is when licensed practitioners offer informal advice outside of a patient-provider context, a phenomenon now regularly witnessed on social media. Claudia Haupt refers to this as "pseudoprofessional advice" and notes that it largely falls outside of the confines of the law [36]. However, she argues that "regulatory interventions can be justified to tie licensed professionals' speech to professional knowledge in order to safeguard against harm." One concern

[7]https://www.whitehouse.gov/ostp/ai-bill-of-rights/

[8]https://developer.chrome.com/en/docs/privacy-sandbox/third-party-cookie-phase-out/; https://www.adexchanger.com/content-studio/how-amazon-ads-is-reshaping-contextual-advertising/

[9]https://code-medical-ethics.ama-assn.org/chapters/patient-physician-relationships

[10]https://www.ftc.gov/business-guidance/resources/health-products-compliance-guidance

is that while licensing *should* ensure medical expertise, in the context of social media, in practice it may be more likely to catch character-related misconduct that may be only somewhat related to quality of care [36].

3.2.3 Product Liability. Product liability law holds businesses responsible for dangerous or defective products. It has been used to regulate myriad products ranging from air bags to baby powder to hot coffee. Product liability has not typically been applied to emerging technologies like social media, advertising, or AI, because it is often not clear what the product is or what the harms are. Courts have typically struck down product liability claims (e.g. Herrick vs. Grindr, where Section 230 protections outweighed the product liability argument [32]). However, that may be changing. There are multiple cases in the past year that have turned to product liability and, in early 2024, a Los Angeles Judge ruled that a Snap case can move to trial [35]. This case was filed on behalf of multiple families whose children had died after taking drugs purchased from dealers they connected with via Snapchat.[11] The youth had believed they were purchasing legitimate drugs which were actually counterfeit, made with fentanyl, and lethal. The lawsuit alleged that Snapchat's design facilitated drug dealers to connect with minors. The crux of the argument by the plaintiffs is that the design not only enables, but encourages, the harmful behavior. The counterargument by the companies is that the problem is the user behavior.

Some of the product liability cases involve severe emotional and sometimes bodily harm, often related to sexual harassment and violence or death. A question in the context of self-diagnosis and targeted ads is whether the harms that might arise from people's consumption of ads rise to the level of damages; that is, should the law step in. So far we do not see evidence of that, at least relative to the kinds of cases in courts today, but often it is hard to see until it is too late (e.g. a tragedy occurs). Also, to apply product liability to the advertisement, there would need to be some kind of misrepresentation or negligence in the advertisement, not just the product [59]. It is conceivable that a minor could see a targeted ad for depression, look for pharmaceuticals to treat it online, and suffer bodily or mental harm. Product liability has some potential for protecting users from being exposed to targeted ads about them, but it may not be invoked in the courts until after a serious harm linked to targeted advertising has taken place.

4 Limitations and Future Research

This study captures self-report data from both the survey and the content analysis. Self-report data may not accurately reflect attitudes or behavior, and may reflect biases towards what participants believe researchers want to see or what social media users believe their audiences want to see. This study was inductive and exploratory, and future work could test how people respond real-time when they see a targeted ad for a sensitive topic in their feeds. It could also measure subsequent behaviors, such as Internet searches for the topic or even physical or behavioral changes (e.g. going to the doctor). We also only measured some possible types of ads that can induce self-diagnosis, there will be others we did not ask about and do not capture (e.g. abusive relationships).

The data in the survey and sampled online reflect just a subset of the U.S. population and other groups in the U.S. and both locally and globally, people might react to targeted ads differently. Some people may also be more prone to self-diagnosis which this study did not measure for. Though we explore potential mechanisms that explain self-diagnosis from targeted ads, we do not provide evidence of proof for those mechanisms. There will of course be other mechanisms that explain people's reactions.

[11] https://www.cagoldberglaw.com/c-a-goldberg-pllc-and-social-media-victims-law-center-file-suit-against-snap-inc-snapchat/

5 Conclusion

People can engage in self-diagnostic reactions to targeted advertisements. This article describes how users react to seeing a sensitive ad topic targeted to them. They may experience anxiety or insecurity, and they make inferences about themselves or other people based on associations implied in a targeted ad. This article considers how repeated exposure to an ad, delivered via an algorithm, might explain the tendency to engage in self-diagnosis. It then considers regulatory protections for self-diagnosis from targeted ads.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grants 1763297 and 2311102.

References

[1] 2019. Assistant Secretary of Fair Housing & Equal Opportunity v. Facebook, Inc. https://www.hud.gov/sites/dfiles/Main/documents/HUD_v_Facebook.pdf

[2] 2017. Do not get sold on drug advertising. Harvard Health Publishing (2017). https://www.health.harvard.edu/medications/do-not-get-sold-on-drug-advertising

[3] Meryl Alper, Jessica Sage Rauchberg, Ellen Simpson, Josh Guberman, and Sarah Feinberg. 2023. TikTok as algorithmically mediated biographical illumination: Autism, self-discovery, and platformed diagnosis on# autisktok. New Media & Society (2023), 14614448231193091.

[4] Tawfiq Ammari, Sarita Schoenebeck, and Meredith Morris. 2014. Accessing social support and overcoming judgment on social media among parents of children with special needs. In Proceedings of the international AAAI conference on web and social media, Vol. 8. 22–31.

[5] Nazanin Andalibi, Oliver L Haimson, Munmun De Choudhury, and Andrea Forte. 2018. Social support, reciprocity, and anonymity in responses to sexual abuse disclosures on social media. ACM Transactions on Computer-Human Interaction (TOCHI) 25, 5 (2018), 1–35.

[6] Nazanin Andalibi, Pinar Ozturk, and Andrea Forte. 2017. Sensitive self-disclosures, responses, and social support on Instagram: the case of# depression. In Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing. 1485–1500.

[7] Theo Araujo, Natali Helberger, Sanne Kruikemeier, and Claes H De Vreese. 2020. In AI we trust? Perceptions about automated decision-making by artificial intelligence. AI & society 35 (2020), 611–623.

[8] Ofer Arazy and Oded Nov. 2010. Determinants of wikipedia quality: the roles of global and local contribution inequality. In Proceedings of the 2010 ACM conference on Computer supported cooperative work. 233–236.

[9] Jack Bandy. 2021. Problematic machine behavior: A systematic literature review of algorithm audits. Proceedings of the acm on human-computer interaction 5, CSCW1 (2021), 1–34.

[10] Kristen Barta and Nazanin Andalibi. 2021. Constructing Authenticity on TikTok: Social Norms and Social Support on the" Fun" Platform. Proceedings of the ACM on Human-Computer Interaction 5, CSCW2 (2021), 1–29.

[11] Aparajita Bhandari and Sara Bimo. 2022. Why's everyone on TikTok now? The algorithmized self and the future of self-making on social media. Social media+ society 8, 1 (2022), 20563051221086241.

[12] Sophie C Boerman, Sanne Kruikemeier, and Frederik J Zuiderveen Borgesius. 2017. Online behavioral advertising: A literature review and research agenda. Journal of advertising 46, 3 (2017), 363–376.

[13] Christina Caron. 2022. Teens Turn to TikTok in Search of a Mental Health Diagnosis. (2022).

[14] Stevie Chancellor, Andrea Hu, and Munmun De Choudhury. 2018. Norms matter: contrasting social support around behavior change in online weight loss communities. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. 1–14.

[15] Stevie Chancellor, Zhiyuan Lin, Erica L Goodman, Stephanie Zerwas, and Munmun De Choudhury. 2016. Quantifying and predicting mental illness severity in online pro-eating disorder communities. In Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing. 1171–1184.

[16] Stevie Chancellor, George Nitzburg, Andrea Hu, Francisco Zampieri, and Munmun De Choudhury. 2019. Discovering alternative treatments for opioid use recovery using social media. In Proceedings of the 2019 CHI conference on human factors in computing systems. 1–15.

[17] Eun Kyoung Choe, Bongshin Lee, Matthew Kay, Wanda Pratt, and Julie A Kientz. 2015. SleepTight: low-burden, self-monitoring technology for capturing and reflecting on sleep behaviors. In Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing. 121–132.

[18] Danielle Keats Citron and Daniel J Solove. 2022. Privacy harms. BUL Rev. 102 (2022), 793.
[19] Kelley Cotter, Julia R DeCook, Shaheen Kanthawala, and Kali Foyle. 2022. In FYP we trust: The divine force of algorithmic conspirituality. International Journal of Communication 16, 2022 (2022), 1–23.
[20] Munmun De Choudhury, Emre Kiciman, Mark Dredze, Glen Coppersmith, and Mrinal Kumar. 2016. Discovering shifts to suicidal ideation from mental health content in social media. In Proceedings of the 2016 CHI conference on human factors in computing systems. 2098–2110.
[21] Nathalie Dierkens. 1991. Information asymmetry and equity issues. Journal of financial and quantitative analysis 26, 2 (1991), 181–199.
[22] Tessa Eagle and Kathryn E Ringland. 2023. "You Can't Possibly Have ADHD": Exploring Validation and Tensions around Diagnosis within Unbounded ADHD Social Media Communities. In Proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility. 1–17.
[23] Daniel A Epstein, Clara Caldeira, Mayara Costa Figueiredo, Xi Lu, Lucas M Silva, Lucretia Williams, Jong Ho Lee, Qingyang Li, Simran Ahuja, Qiuer Chen, et al. 2020. Mapping and taking stock of the personal informatics literature. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 4, 4 (2020), 1–38.
[24] Alexandre Faisal-Cury, Carolina Ziebold, Daniel Maurício de Oliveira Rodrigues, and Alicia Matijasevich. 2022. Depression underdiagnosis: Prevalence and associated factors. A population-based study. Journal of psychiatric research 151 (2022), 157–165.
[25] Casey Fiesler and Nicholas Proferes. 2018. "Participant" perceptions of Twitter research ethics. Social Media+ Society 4, 1 (2018), 2056305118763366.
[26] Lucy Foulkes and Jack L Andrews. 2023. Are mental health awareness efforts contributing to the rise in reported mental health problems? A call to test the prevalence inflation hypothesis. New Ideas in Psychology 69 (2023), 101010.
[27] Amy Gaeta. 2023. Diagnostic advertisements: The phantom disabilities created by social media surveillance. (2023).
[28] Liza Gak, Seyi Olojo, and Niloufar Salehi. 2022. The distressing ads that persist: Uncovering the harms of targeted weight-loss ads among users with histories of disordered eating. Proceedings of the ACM on Human-Computer Interaction 6, CSCW2 (2022), 1–23.
[29] Doğa Gatos, Aslı Günay, Güncel Kırlangıç, Kemal Kuscu, and Asim Evren Yantac. 2021. How HCI bridges health and design in online health communities: a systematic review. In Designing Interactive Systems Conference 2021. 970–983.
[30] John Gilmour, Tanya Machin, Charlotte Brownlow, and Carla Jeffries. 2020. Facebook-based social support and health: A systematic review. Psychology of popular media 9, 3 (2020), 328.
[31] Ylva Ginsberg, Javier Quintero, Ernie Anand, Marta Casillas, and Himanshu P Upadhyaya. 2014. Underdiagnosis of attention-deficit/hyperactivity disorder in adult patients: a review of the literature. The primary care companion for CNS disorders 16, 3 (2014), 23591.
[32] Carrie Goldberg. 2020. Winning Through Losing. Perspectives 28 (2020), 7.
[33] Cami Goray and Sarita Schoenebeck. 2022. Youths' Perceptions of Data Collection in Online Advertising and Social Media. Proceedings of the ACM on Human-Computer Interaction 6, CSCW2 (2022), 1–27.
[34] Chang-Dae Ham. 2017. Exploring how consumers cope with online behavioral advertising. International Journal of Advertising 36, 4 (2017), 632–658.
[35] Taylor Hatmaker. 2024. Lawsuit against Snap over fentanyl deaths can proceed, judge rules. (2024).
[36] Claudia E Haupt. 2022. Pseudoprofessional Advice. Boston University Law Review 424 (2022).
[37] Claudia E Haupt and Mason Marks. 2023. AI-generated medical advice—GPT and beyond. Jama 329, 16 (2023), 1349–1350.
[38] Kashmir Hill. 2012. How target figured out a teen girl was pregnant before her father did. Forbes, Inc 7 (2012), 4–1.
[39] Claudiu Gabriel Ionescu and Monica Licu. 2023. Are TikTok algorithms influencing users' self-perceived identities and personal values? A mini review. Social Sciences 12, 8 (2023), 465.
[40] Annemarie Jutel. 2017. "Dr. Google" and his predecessors. Diagnosis 4, 2 (2017), 87–91.
[41] Nadia Karizat, Dan Delmonaco, Motahhare Eslami, and Nazanin Andalibi. 2021. Algorithmic folk theories and identity: How TikTok users co-produce Knowledge of identity and engage in algorithmic resistance. Proceedings of the ACM on human-computer interaction 5, CSCW2 (2021), 1–44.
[42] Jon Keegan and Joel Eastwood. 2023. From "Heavy Purchasers" of Pregnancy Tests to the Depression-Prone: We Found 650,000 Ways Advertisers Label You. (2023).
[43] Jeff Kosseff. 2022. A User's Guide to Section 230, and a Legislator's Guide to Amending It (or Not). Berkeley Tech. LJ 37 (2022), 757.
[44] Melvin Kranzberg. 1986. Technology and history:" Kranzberg's laws". Technology and culture 27, 3 (1986), 544–560.
[45] Jennifer L Kraschnewski, Cynthia H Chuang, Erika S Poole, Tamara Peyton, Ian Blubaugh, Jaimey Pauli, Alyssa Feher, and Madhu Reddy. 2014. Paging "Dr. Google": does technology fill the gap created by the prenatal care visit structure? Qualitative focus group study with pregnant women. Journal of medical Internet research 16, 6 (2014), e3385.

[46] Srijan Kumar, Robert West, and Jure Leskovec. 2016. Disinformation on the web: Impact, characteristics, and detection of wikipedia hoaxes. In Proceedings of the 25th international conference on World Wide Web. 591–602.
[47] Michelle S Lam, Ayush Pandit, Colin H Kalicki, Rachit Gupta, Poonam Sahoo, and Danaë Metaxa. 2023. Sociotechnical Audits: Broadening the Algorithm Auditing Lens to Investigate Targeted Advertising. Proceedings of the ACM on Human-Computer Interaction 7, CSCW2 (2023), 1–37.
[48] Angela Y Lee, Hannah Mieczkowski, Nicole B Ellison, and Jeffrey T Hancock. 2022. The algorithmic crystal: Conceptualizing the self through algorithmic personalization on TikTok. Proceedings of the ACM on Human-computer Interaction 6, CSCW2 (2022), 1–22.
[49] Kenneth Lee, Kreshnik Hoti, Jeffery David Hughes, and Lynne M Emmerton. 2015. Consumer use of "Dr Google": a survey on health information-seeking behaviors and navigational needs. Journal of medical Internet research 17, 12 (2015), e288.
[50] Gregory Lewis. 2011. Asymmetric information, adverse selection and online disclosure: The case of eBay motors. American Economic Review 101, 4 (2011), 1535–1546.
[51] Leslie S Liu, Jina Huh, Tina Neogi, Kori Inkpen, and Wanda Pratt. 2013. Health vlogger-viewer interaction in chronic illness management. In Proceedings of the SIGCHI conference on Human factors in computing systems. 49–58.
[52] Haiwei Ma, C Estelle Smith, Lu He, Saumik Narayanan, Robert A Giaquinto, Roni Evans, Linda Hanson, and Svetlana Yarosh. 2017. Write for life: Persisting in online health communities through expressive writing and social support. Proceedings of the ACM on Human-Computer Interaction 1, CSCW (2017), 1–24.
[53] Catharine A MacKinnon. 2020. Weaponizing the First Amendment. Virginia Law Review 106, 6 (2020), 1223–1283.
[54] Colin MacLeod and Lynlee Campbell. 1992. Memory accessibility and probability judgments: an experimental evaluation of the availability heuristic. Journal of personality and social psychology 63, 6 (1992), 890.
[55] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. Proceedings of the ACM on human-computer interaction 3, CSCW (2019), 1–23.
[56] Matthew B Miles and A Michael Huberman. 1994. Qualitative data analysis: An expanded sourcebook. sage.
[57] Michael L Millenson, Jessica L Baldwin, Lorri Zipperer, and Hardeep Singh. 2018. Beyond Dr. Google: the evidence on consumer-facing digital tools for diagnosis. Diagnosis 5, 3 (2018), 95–105.
[58] Ashlee Milton, Leah Ajmani, Michael Ann DeVito, and Stevie Chancellor. 2023. "I See Me Here": Mental Health Content, Community, and Algorithmic Curation on TikTok. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. 1–17.
[59] Fred W Morgan and Jeffrey J Stoltman. 1997. Advertising and product liability litigation. Journal of Advertising 26, 2 (1997), 63–75.
[60] Albert G Mulley, Chris Trimble, and Glyn Elwyn. 2012. Stop the silent misdiagnosis: patients' preferences matter. Bmj 345 (2012).
[61] David E Newman-Toker, Najlla Nassery, Adam C Schaffer, Chihwen Winnie Yu-Moe, Gwendolyn D Clemens, Zheyu Wang, Yuxin Zhu, Ali S Saber Tehrani, Mehdi Fanai, Ahmed Hassoon, et al. 2023. Burden of serious harms from diagnostic error in the USA. BMJ Quality & Safety (2023).
[62] Fayika Farhat Nova, Amanda Coupe, Elizabeth D Mynatt, Shion Guha, and Jessica A Pater. 2022. Cultivating the community: inferring influence within eating disorder networks on twitter. Proceedings of the ACM on Human-Computer Interaction 6, GROUP (2022), 1–33.
[63] Paul Ohm. 2014. Sensitive information. S. Cal. L. Rev. 88 (2014), 1125.
[64] Scott B Patten. 2012. Prevalence Inflation. 649–650 pages.
[65] Elissa M Redmiles, Jessica Bodford, and Lindsay Blackwell. 2019. "I just want to feel safe": A Diary Study of Safety Perceptions on Social Media. In Proceedings of the International AAAI Conference on Web and Social Media, Vol. 13. 405–416.
[66] Howard Rheingold. 2000. The virtual community, revised edition: Homesteading on the electronic frontier. MIT press.
[67] Princess Sampson, Ro Encarnacion, and Danaë Metaxa. 2023. Representation, Self-Determination, and Refusal: Queer People's Experiences with Targeted Advertising. In Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency. 1711–1722.
[68] Sarita Schoenebeck, Cami Goray, Amulya Vadapalli, and Nazanin Andalibi. 2023. Sensitive Inferences in Targeted Advertising. Nw. J. Tech. & Intell. Prop. 21 (2023), 155.
[69] Norbert Schwarz, Herbert Bless, Fritz Strack, Gisela Klumpp, Helga Rittenauer-Schatka, and Annette Simons. 1991. Ease of retrieval as information: Another look at the availability heuristic. Journal of Personality and Social psychology 61, 2 (1991), 195.
[70] Daniel J Solove. 2012. Introduction: Privacy self-management and the consent dilemma. Harv. L. Rev. 126 (2012), 1880.

[71] Daniel J Solove. 2023. Data is what data does: Regulating Use, Harm, and risk instead of sensitive data. Harm, and Risk Instead of Sensitive Data (January 11, 2023) (2023).
[72] Alicia Solow-Niederman. 2022. Information privacy and the inference economy. Nw. UL Rev. 117 (2022), 357.
[73] Shaina J Sowles, Monique McLeary, Allison Optican, Elizabeth Cahn, Melissa J Krauss, Ellen E Fitzsimmons-Craft, Denise E Wilfley, and Patricia A Cavazos-Rehg. 2018. A content analysis of an online pro-eating disorder community on Reddit. Body image 24 (2018), 137–144.
[74] David R Stukus. 2019. How Dr Google is impacting parental medical decision making. Immunology and Allergy Clinics 39, 4 (2019), 583–591.
[75] Hyewon Suh, Nina Shahriaree, Eric B Hekler, and Julie A Kientz. 2016. Developing and validating the user burden scale: A tool for assessing user burden in computing systems. In Proceedings of the 2016 CHI conference on human factors in computing systems. 3988–3999.
[76] Briony Swire-Thompson, David Lazer, et al. 2020. Public health and online misinformation: challenges and recommendations. Annu Rev Public Health 41, 1 (2020), 433–451.
[77] David R Thomas. 2006. A general inductive approach for analyzing qualitative evaluation data. American journal of evaluation 27, 2 (2006), 237–246.
[78] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In proceedings of the eighth symposium on usable privacy and security. 1–15.
[79] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un) informed consent: Studying GDPR consent notices in the field. In Proceedings of the 2019 acm sigsac conference on computer and communications security. 973–990.
[80] Jojanneke Van der Toorn, Tom R Tyler, and John T Jost. 2011. More than fair: Outcome dependence, system justification, and the perceived legitimacy of authority figures. Journal of experimental social psychology 47, 1 (2011), 127–138.
[81] Kaan Varnali. 2021. Online behavioral advertising: An integrative review. Journal of Marketing Communications 27, 1 (2021), 93–114.
[82] Salome Viljoen. 2021. A relational theory of data governance. Yale LJ 131 (2021), 573.
[83] Sandra Wachter and Brent Mittelstadt. 2019. A right to reasonable inferences: re-thinking data protection law in the age of big data and AI. Colum. Bus. L. Rev. (2019), 494.
[84] Ryen W White and Eric Horvitz. 2009. Cyberchondria: studies of the escalation of medical concerns in web search. ACM Transactions on Information Systems (TOIS) 27, 4 (2009), 1–37.
[85] Ryen W White and Eric Horvitz. 2013. Captions and biases in diagnostic search. ACM Transactions on the Web (TWEB) 7, 4 (2013), 1–28.
[86] Ryen W White and Eric Horvitz. 2015. Belief dynamics and biases in web search. ACM Transactions on Information Systems (TOIS) 33, 4 (2015), 1–46.
[87] Yuxi Wu, Sydney Bice, W Keith Edwards, and Sauvik Das. 2023. The Slow Violence of Surveillance Capitalism: How Online Behavioral Advertising Harms People. In Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency. 1826–1837.
[88] Yaxing Yao, Davide Lo Re, and Yang Wang. 2017. Folk models of online behavioral advertising. In Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing. 1957–1969.

Received January 2024; revised July 2024; accepted October 2024