Intuitions, Analytics, and Killing Ants: Inference Literacy of High School-educated Adults in the US

Jeffrey Warshaw, *University of California, Santa Cruz;*
Nina Taft and Allison Woodruff, *Google, Inc.*

https://www.usenix.org/conference/soups2016/technical-sessions/presentation/warshaw

This paper is included in the Proceedings of the
Twelfth Symposium on Usable Privacy and Security (SOUPS 2016).

June 22–24, 2016 • Denver, CO, USA

ISBN 978-1-931971-31-7

Open access to the Proceedings of the
Twelfth Symposium on Usable Privacy
and Security (SOUPS 2016)
is sponsored by USENIX.

Intuitions, Analytics, and Killing Ants: Inference Literacy of High School-educated Adults in the US

Jeffrey Warshaw*
University of California, Santa Cruz
1156 High Street
Santa Cruz, CA, USA 95064
jwarshaw@ucsc.edu

Nina Taft
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
ninataft@google.com

Allison Woodruff
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
woodruff@acm.org

ABSTRACT

Analytic systems increasingly allow companies to draw inferences about users' characteristics, yet users may not fully understand these systems due to their complex and often unintuitive nature. In this paper, we investigate inference literacy: the beliefs and misconceptions people have about how companies collect and make inferences from their data. We interviewed 21 non-student participants with a high school education, finding that few believed companies can make the type of deeply personal inferences that companies now routinely make through machine learning. Instead, most participant's inference literacy beliefs clustered around one of two main concepts: one cluster believed companies make inferences about a person based largely on a priori stereotyping, using directly gathered demographic data; the other cluster believed that companies make inferences based on computer processing of online behavioral data, but often expected these inferences to be limited to straightforward intuitions. We also find evidence that cultural models related to income and ethnicity influence the assumptions that users make about their own role in the data economy. We share implications for research, design, and policy on tech savviness, digital inequality, and potential inference literacy interventions.

1. INTRODUCTION

The ways that companies gain insights from consumer data have changed drastically in the last few decades, and yet we know little about how the general public's understanding has kept up with those changes. Many decisions that companies historically made through market research and coarse, demographic segmentation are now instead driven by statistical inferencing, through online data mining and machine learning. The ability to algorithmically process behavioral data and look for patterns across millions of users allows companies to infer characteristics that users may believe are difficult to guess or hidden online, such as their hobbies and likes [35], their age and ethnicity [30,51], and their personality and values [17,18]. These inferences are used in a wide range of everyday contexts, for example to offer personalized ads and product recommendations, or to offer differentiated pricing or employment opportunities [14,43].

Because algorithmic inferences can have economic and other far-reaching consequences in people's lives [14,43], it can be valuable for people to have an understanding of what can be inferred about them and how. However, the systems that generate these inferences are often complex and/or opaque. Recent research has emphasized the surprise that many users experience when learning about inferential systems [18,58,61], implying a gap likely exists between what people generally believe companies currently do with their data, and what the state-of-the-art actually is. To date, though, research on *digital literacy* has focused on knowledge of data collection practices [5,44,57] but to our knowledge has not explored beliefs and misconceptions people hold about companies' inferencing methods and capabilities. We argue for the inclusion of these beliefs as a subconstruct of digital literacy, and we introduce the term *inference literacy* to describe it.

In this work, we share results from a qualitative study assessing the inference literacy of 21 US non-student adults with a high school degree but no post-secondary degree, the modal educational attainment in the US, comprising 49% of the adult US population [60]. Inspired by previous work on folk models [62,63], we explored beliefs and misconceptions, and found two distinct clusters within our sample. One cluster believed that online companies rely on now-outdated market research strategies that companies used decades ago [22], such as data collection through surveys rather than through tracking user behavior online. This cluster also interpreted inferential techniques used by companies as constituting *stereotyping*, and expressed worries about hackers and scammers. The other cluster believed that companies mine people's online behavior to infer their preferences, using computer analytics to make intuitive predictions about users. Neither cluster had fully accurate beliefs, and both clusters had misconceptions that have important user experience implications.

Further, we argue for the broadening of cross-cultural studies in usable privacy and security to explicitly include qualitative differences based on social class and ethnicity rather than just on national culture. Building on research that has examined folk models that people have about online phenomena at an individual [62] or national [28] level, we provide evidence within our high school-educated sample that users' interpretation of the privacy ecosystem can vary substantially based on social class and

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

Symposium on Usable Privacy and Security (SOUPS) 2016, June 22-24, 2016, Denver, Colorado.

* Work performed while at Google Inc.

ethnicity. We link this to cognitive anthropology research on *cultural models*, the sets of assumptions that members of a group form over time based on shared experiences [11,23]. Cultural models about personal agency and choice, both of which might affect a person's online privacy beliefs and behavior, vary between socioeconomic and ethnic groups [54,56,23]. In this work, we saw differences in the framing of privacy decisions as risks or choices as a function of participant income, relating to differences in cultural models of personal agency based on social class [54,56,23]. We also found that ethnic minority participants interpreted companies inferring their preferences based on their ethnicity as stereotyping, which we contextualize in terms of social psychology research on ethnic minority groups' experiences with discrimination in consumer settings [15,34,50].

Our main contributions are:

- We present a novel study of *inference literacy*, describing the beliefs and misconceptions that 21 US adults with a high school education hold about how companies make inferences from their online data.
- We report and describe two clusters of beliefs that together describe 19 out of our 21 participants. These clusters link inference literacy to cultural models of agency and point to apparent digital inequalities based on socioeconomic factors, including income and ethnicity.
- We argue for the redefinition of tech savviness and digital literacy to include inference literacy, as well as for cultural models based on social class and ethnicity to be included in future online privacy research.

2. RELATED WORK

2.1 Inferencing methods

Our work explores users' understanding of current inferencing methods commonly used by companies. Companies currently rely heavily on machine learning to make inferences about users, applying techniques such as supervised or unsupervised learning, reinforcement learning, deep learning, or neural networks to find complex patterns in behavioral data [2,4,41]. These methods are typically applied to datasets containing multiple streams of data aggregated from large groups of users in order to find correlations between variables of interest [6]. These techniques can uncover strong correlations, similar to those humans might intuitively guess when presented with frequently co-occurring behaviors; they can also uncover weaker, more unintuitive correlations that are only detectable by combining data from a large number of users. A key property of the relationships between variables that these techniques uncover is that they are generalizable to new users: learning how to predict variable A from variables B, C, and D based on a group of people who shared all of those variables enables the creation of a system that can infer variable A for people who have not shared it. For example, analyses of large datasets have led to systems that can infer a person's gender based on their movie ratings [66], their religion from their search queries [7], their sexual preference from their social media likes [30], and their personality and values from their social media text [17,18]. A system's confidence in inferring an unknown characteristic typically increases with the number of predictor variables available, but even a small number of data points can be used to make a better-than-random guess, e.g., [30].

Despite the key role that machine learning systems play in the data economy, the workings of inferencing systems are often opaque, lacking transparency to users about what data they use or how they work. There are a few efforts that have studied users' reactions to inferencing systems. [61] presented social media users with personality profiles that an inferential system automatically generated from their social media posts, and users' reactions spanned from surprise at how accurate the profiles were, to creepiness and learned helplessness about whether they could decline to share them in different settings. In [58], the authors studied users' reactions to online behavioral advertising and found that they felt it was both useful and scary at the same time. Kulesza and colleagues found that having the ability to correct an automatic recommender system does not in itself improve users' mental models of the process by which the system makes inferences [31,32]. Instead, they saw participants' confidence in unsound mental models increased over time unless they received a structured educational intervention prior to using the system [31]. In the current study, we assess people's global beliefs about what data companies use in inferencing, what methods companies use to make inferences, and limitations of inferencing systems.

2.2 Folk models of online privacy and security

Several recent online privacy and security studies have explored folk models, sets of beliefs and misconceptions that non-experts have about a particular topic. Rather than assuming non-experts have zero knowledge, folk models acknowledge that people develop their own lay theories to explain ambiguous situations they encounter. The online privacy landscape is often ambiguous, leaving users to come up with their own incomplete explanations for phenomena like hackers and viruses [28,62]. Research on users' understanding about online data collection and specific inferential systems has found non-experts often have consequential misunderstandings about the online landscape [26], including systems they commonly interact with such as autocomplete [48] or behavioral advertising [39,58]. We build on and extend this previous work by exploring folk models related to inferencing.

Importantly, folk models are not independently produced. Social and cultural factors affect them as well. Informal stories and advice about privacy and security are commonly shared [49]. These informal sources of information may include out-of-date information, with some security advice that non-experts endorse being decades behind what experts recommend based on current threats [25]. This *social* aspect of folk models has been discussed in research in online privacy and security, but *cultural* influences are rarely discussed. Folk models of viruses and hackers do appear to differ cross-culturally [28,62], but neither study examined the relationship between culture and the different folk models evidenced in their samples. In the current study, we explore folk models of inference literacy through an explicitly cultural lens, examining the role of cultural models in shaping beliefs and attitudes about the data economy.

2.3 Cultural models and technology

Cultural models are sets of implicit assumptions that develop based on shared experiences and common history, which differ qualitatively between different cultural groups [11,23]. Culture has often been applied in HCI to describe differences based on national culture or language [9,13,21], but other fields describe substantial differences in cultural models based on other features, including educational attainment, social class, and ethnicity [47,54,56]. Recent research has described cross-cultural differences in definitions of privacy based on national or religious culture, and how those differences relate to the trade-offs users make in online settings [1,59]. In the current work, we link inference literacy beliefs and attitudes to two specific types of cultural models. First, we draw on research showing different cultural models of personal agency for middle-class and working-class Americans [54,56]. Middle-class Americans typically develop a more *independent* sense of agency, expecting to exercise control within their environment. By comparison, working-class Americans tend to develop a more *interdependent* sense of agency, expecting to cope with external factors rather than exercising independent choice themselves. This stems from differences in economic and environmental constraints between these groups [56], and is reinforced by socialization and media consumption that differ by education and income [40,54].

We also draw on research on experiences of marginalized ethnic groups with stereotyping and discrimination. Ethnic minorities in the US often encounter stereotyping and discrimination across various settings, including in education [3,55], while driving [37], and as pedestrians [16]. This pattern extends to consumer experiences as well. Research on "shopping while Black" has shown substantial differences in customers' treatment in US retail settings based on their ethnicity [34,50,52]. Ambiguously discriminatory experiences such as being ignored, followed by retail staff, or not given service can be interpreted as an institutional distrust for or devaluing of them based on their ethnicity [34,50,52]. Inferential systems can themselves be ambiguous to users in how they operate, and the line between personalization and stereotyping may not always be clear. In this study, we include participants from marginalized ethnic groups to obtain their perspectives on this and other inferencing topics.

2.4 Digital inequality

Several studies have looked at digital inequality: ways that offline socioeconomic inequalities related to demographic categories like educational attainment, income, ethnicity, and age are reproduced in online settings. There are differences in internet usage by educational attainment. Those with a high school degree tend to engage in fewer different types of online activities [46,65], fewer capital-enhancing activities online [20,46], and are more likely to be reliant on a smartphone rather than computer to access the internet than the college-educated [53]. Despite the common belief that the "do-it-yourself" opportunities that online access enables are sufficient to decrease social inequalities, socioeconomically disadvantaged internet users benefit from online access at the same or slower rates compared to those with higher incomes or education [12,38,45]. High impact decisions of inferential systems such as credit scoring often contain systematic errors or design decisions that disproportionately disadvantage those with lower SES or non-European-American ethnicity [14,43]. In the current study, we explore a potential digital inequality: whether differences in inference literacy are related to socioeconomic status.

3. METHOD

For our study of inference literacy, we collected data from a sample of US adults with a high school education to learn their beliefs and misconceptions about companies' inferencing methods. Each session contained two main sections that took place consecutively: an interview to elicit the participant's existing beliefs about what, how, and why companies collect and use their data; and a teaching intervention for the participant to learn two basic principles of current inferencing methods. We focus in this paper mainly on data from the interview section, but we include relevant details about the design of the teaching section in Appendix A.

3.1 Study Design

To explore inference literacy in our participants, we adapted Oakleaf's "Information Literacy Instruction Assessment Cycle" (ILIAC) [42]. The ILIAC is an iterative educational research method that aids in creating learning activities that assess a student's knowledge before, during, and after the activity. This method has been applied successfully in educational settings where the goal is to assess and teach digital literacy concepts in a single session [42].

We adapted this method to fit the current study, going through four full cycles, each of which took 1-3 weeks and included 2-12 participants. The overwhelming majority of changes were made to the teaching procedure, with only minor wording changes made to interview questions between cycles.

3.2 Participants

We collected data from 23 participants in total between July and September of 2015, all of whom were recruited by a research recruitment firm with a respondent database containing San Francisco Bay Area residents. We recruited participants who had a high school degree or the equivalent (i.e., GED) but no post-secondary degree, and who were not currently enrolled in post-secondary education. In addition, we aimed to explore socioeconomic and cultural differences in folk models of inference literacy, so we recruited a diverse sample in terms of age, gender, ethnicity, household income, parental status, occupation, and political beliefs. We created a recruitment screener that asked about these demographic categories, as well as several other questions, such as which internet-accessible devices the participant owned, and news sources the participant uses.

Two participants out of the 23 participated in a pilot study. Because the procedure changed significantly based on the pilot, we exclude these pilot participants from the analysis reported here. The 21 participants in the final sample include 10 women and 11 men, ranging in age from 18 to over 65 years of age. Eleven participants identified as White or European-American, 5 as Black or African-American, 3 as Asian-American or Pacific Islander, 3 as Hispanic or Latino, and 3 identified as having mixed or multiple ethnicities. Occupations were varied, including waste management driver, payroll clerk, security guard, HVAC technician, retired, and unemployed. Participants were interviewed in-person in one of two locations, Mountain View, California (n=16), or San Francisco, California (n=5), and were compensated for their time.

3.3 Session Procedure

We first describe the general structure of the session procedure and then detail each component in the order participants experienced it. The same interviewer led each participant through a 90-minute session with two main components: (a) a semi-structured interview meant to elicit existing beliefs and misconceptions about how companies make inferences from their data, and (b) a teaching intervention during which participants engaged with real world examples of inferences that companies can or cannot make from data. Prior to each session, the interviewer verbally walked the participant through an informed consent form that described the study. With participants' permission, each session was video recorded to facilitate transcription and coding. Another member of the research team observed each session from a separate room either during or after the session, taking notes that included quotes representative of that participant's beliefs, and preliminary themes that arose across multiple participants' sessions. Between sessions, the research team frequently met to discuss observations, develop the analysis plan, and make changes to session procedure for future cycles.

3.3.1 Belief elicitation interview

Inspired by Wash's work on eliciting participants' lay beliefs about home security [62], the first portion of each session consisted of a semi-structured interview that we developed to learn participants' existing beliefs about how companies collect their data and use it to make inferences about them. The interviewer used a paper script containing questions to facilitate the conversation, and began by asking participants their educational background, occupation, and familiarity with machine learning. Only two participants had heard of machine learning, both of whom claimed it referred to some kind of rudimentary artificial intelligence.

To ground the belief elicitation interview in terms of each participant's daily experiences, the first and main prompt for each participant was, "Think about what you'll do online today, and talk me through things that companies will try to figure out about you based on what you do online today". The responses to this prompt were detailed and varied. Participants referred to different settings, with some referring exclusively to smartphones or laptops whereas others described mixed usage of devices. We did not constrain the companies participants talked about, and they described interactions with a wide range of companies for a variety of tasks, including checking email, social media browsing and posting, online banking, retail browsing and purchasing, and watching videos online. Many beliefs about how companies collect and use data came out naturally as participants described their daily online experiences. If they did not arise spontaneously during the interview, the interviewer asked follow-up questions to elicit more detail on each participant's beliefs, including whether, why, and how they believe companies collect data; what kinds of data companies do and do not collect; and whether and how companies make guesses about individuals' characteristics. After probing the contents and sources for each of these potential beliefs, the interviewer concluded the belief elicitation interview and moved to the teaching intervention phase of the session.

3.3.2 Teaching intervention

The goal of the teaching intervention section was to assess participants' explanations about the inferencing processes and capabilities companies deploy, before and after providing participants with brief explanations about modern data collection and inferencing phenomena. After each explanation, the interviewer conducted a card-sorting task with the participant where they rated and discussed the likelihood that companies can make a particular inference from a particular type of user data. Because the focus of the current paper is on participants' pre-existing beliefs, much of the data collected in this section is outside the scope of the coding and results described in this paper. We did use participants' responses to the pre-test assessment, as they were directly relevant to beliefs about inferencing, and the pre-test was given prior to any teaching: "If a social media company wants to learn more about their users, what would they be able to figure out about a user even if that person didn't tell them? How would they figure that out? What would be impossible for a social media company to figure out about someone?" We also used a small number of beliefs that participants shared after the teaching intervention where it was clear that these were pre-existing, e.g., "I always thought it was X" after we taught them Y. We include the remainder of the teaching intervention procedure as Appendix A.

3.4 Coding

In this section, we detail the affinity diagramming and coding of participants videos and transcripts that allowed us to characterize participants' inference literacy beliefs and attitudes about the data economy.

We began analyzing the interview data by creating affinity diagrams [19], taxonomies where participants' perspectives could be grouped across various axes. Some of these diagrams were digital, containing quotes from interviews that we sorted according to thematic differences in how participants described inferencing phenomena. Other diagrams were physical, and used the participants as the unit of analysis. These holistic groupings allowed us to tease apart the key components of qualitatively different folk models about data collection and inferencing, as well as to analyze for cultural and socioeconomic themes such as stereotyping and risk perception. The research team discussed these diagrams as they were created, iterating on them several times during the analysis process.

Additionally, we reviewed the transcripts to identify and define codes similar to [8] to describe the wide range of beliefs participants expressed. The interviewer first coded each transcript, obtaining feedback from the entire research team about ambiguous codes. This coding process was iterative, so that transcripts read early on were reviewed to check for codes that were discovered or refined later in the coding process. To establish intercoder reliability [33], another author coded each of the transcripts for the key beliefs described in the results below. Intercoder agreement was above 75% for the first five transcripts analyzed, and above 80% for the first pass through all 21 transcripts. Disagreements between the first and second coder were resolved by reviewing the transcripts and discussing to come to agreement. In the majority of these disagreements, the two coders agreed about the participant's belief but had different opinions about the level of proof required to confidently assign a code. We took a conservative stance in these cases, requiring supporting statements that were unambiguous or repeated during the interview. After revising the codebook and assessing the remaining disagreements, intercoder agreement was above 90%, indicating that the codes were sufficiently well-defined and reliably assigned during the coding process. The final codebook contained 160 unique codes from the 21 participants.

3.5 Clustering

During data collection and coding, we noticed that some beliefs appeared to frequently co-occur and decided to explore this possibility systematically. As our interest was in describing inference literacy, we focused primarily here on beliefs about data collection and inferencing. After coding the transcripts, we collected the 31 inference-related codes that we had assigned to four or more participants. We then created a vector for each code, each containing the list of participants who had been assigned that code. We manually compared the vectors pairwise, looking for

Table 1. Categorized list of codes contained within each inference literacy cluster, with beliefs that formed the initial core of each cluster in bold. Paper sections discussing each code and related results from affinity diagramming are in parentheses.

	Market Research Cluster (n = 8)	Data Mining Cluster (n = 11)
Data collection beliefs	**Companies collect demographics by surveys. (4.1.1)** **Companies collect personal information from public records. (4.1.1)**	**Companies collect online behavioral data. (4.2.1)** Companies doing retail retargeting taught me that my behavior is collected. (4.2.1)
Inferencing beliefs	**Companies make inferences by having humans make common sense intuitions. (4.1.2)** Companies stereotype users based on their demographics. (4.1.3)	Companies make recommendations using behavioral data. (4.2.2) **Companies use computer analytics to make inferences. (4.2.3)** Companies tailor ads based on what you click. (4.2.3) Inferences are made by analyzing your social network. (4.2.3)
Attitudes	Companies stereotyping is morally wrong. (4.1.3) I am worried about hackers. (4.1.4) I am worried about scammers. (4.1.4)	I feel “watched” or “tracked”. (4.2.1)

frequently co-occurring beliefs as well as beliefs that were strongly negatively correlated, such that two vectors had few or no overlapping participants between them.

There were unmistakable links between beliefs about data collection methods and beliefs about inferencing methods that formed the basis for the rest of the clustering process. First, beliefs that online companies collect demographic data by survey or collect personal information by public records were associated with the belief that companies make inferences by relying on common sense intuition rather than computer processing of data. Second, the belief that companies collect online behavioral data overlapped completely with the belief that companies use computer analytics to make inferences. These two sets of beliefs are conceivably complementary in that they each describe one aspect of current inferencing methods, but surprisingly, there was no overlap between these sets of beliefs. Participants who believed that companies use computer analytics did not express that they believed companies collect demographic data by survey, and so on. These two sets of highly distinguishable inference literacy beliefs therefore formed the core for us to explore other connections between our data.

We compared the remaining arrays of belief codes to identify other commonalities, finding two distinct sets of 7 codes that anchored around the core beliefs above. These two clusters of beliefs and attitudes appear in Table 1. Although we began the clustering process seeking to identify sets of beliefs and we did not presuppose that these would be largely mutually exclusive, we found that participants with beliefs in one cluster had few or no beliefs from the other cluster. Because the interviews often surfaced issues related to social class and ethnicity, we holistically analyzed the clusters, drawing on research on cultural models to interpret the codes in light of participant demographics in the results below. Out of 21 participants, 19 were assigned to one of the two clusters. The remaining two participants believed that companies could not or would not collect data about individual users. Although this is a crucial misconception, it was so infrequent that we were unable to explore it systematically in the present study.

The alert reader may wonder whether these clusters constitute “folk models” as described in other literature [28,62]. In that our clusters describe non-expert sets of beliefs held by our participants, it would be reasonable to refer to the clusters as folk models. In this work, we use the word “cluster” for consistency, as it applies equally well to the sets of beliefs themselves and the participants who held them.

3.6 Limitations

We note several limitations of our study methodology that should be considered when interpreting this work. First, due to our focus on describing beliefs of high school-educated adults, we did not include college-educated adults in our sample. This prevents us from comparing inference literacy beliefs across different levels of educational attainment. Second, our sample was not statistically representative of the US adult high school-educated population. The clusters we report should be viewed as a deep exploration of our sample’s beliefs and attitudes, but not as generalizing to that population as a whole. Third, we report several misconceptions that people have about inferencing methods, but we do not have data to say that these misconceptions lead to harmful privacy behaviors. Useful behaviors can arise from incomplete or incorrect beliefs [62,63], and that may be the case here as well. Finally, because we touch on socioeconomic status and ethnicity in this work, we include the detail that the research team consisted only of college-educated, European-American researchers. We describe participants’ experiences in their own words, but our interpretations may lack context or nuance that may have been more readily available to a more diverse research team.

4. RESULTS

Based on our analysis of the interview data, we identified two main clusters of inferencing beliefs held by participants in our sample. The “market research” (MR) cluster was anchored by a shared belief that companies ask users direct questions about their demographics and personally identifiable information, to sell to them based largely on stereotype. The “data mining” (DM) cluster relied on a shared belief that companies track users’ online behavior, to make retail or media recommendations based on their past behavior. We also observed that ethnicity and socioeconomic status were associated with differences in the interpretations participants made about inferencing processes and their own personal agency in the data economy.

Several participants across both clusters had important misconceptions. Participants in both clusters claimed companies rely on human employees to make inferences about users, which we refer to throughout as “humans-in-the-loop”. Related to that misconception, most participants felt inferences are made only based on strong, intuitive connections between two pieces of data, rather than by using multiple pieces of evidence to support an inference. Although data collection practices and inferencing methods differ across companies, our participants rarely made such distinctions.

In the following, we refer to participants with the Market Research cluster of beliefs as MR1, MR2,... MR8, and those with the Data Mining cluster of beliefs as DM1, DM2,... DM11.

4.1 Market Research Cluster

Participants in the MR cluster believed that companies primarily collect data by asking users directly for their personal information (4.1.1), and that companies make shallow, potentially harmful assumptions about them based on their demographics (4.1.2). These participants believed companies make inferences about users based on a priori assumptions about links between two pieces of data, saying things like, "it goes hand-in-hand" (MR1) and "You can make certain summations just by looking at somebody" (MR6). They often described this inferencing process as *stereotyping* (4.1.3), claiming that companies use demographic information like income or ethnicity to make marketing decisions. This was not seen as a benign form of personalization; rather, participants expressed strong moral objections to companies stereotyping in this manner. We also found that despite the interview focusing on companies, participants with these beliefs spontaneously expressed strong concerns about hackers or scammers getting access to their data (4.1.4). This concern about being targeted by criminals often drowned out any apprehension they might otherwise have about what companies would do with their data. In this section, we describe the beliefs belonging to this cluster in greater detail.

4.1.1 Companies collect demographics and personal information from direct sources

The core belief held by participants in the MR cluster was that online companies collect users' demographics and personal information explicitly. There were two main ways they described companies collecting those data: asking a user for it directly in a survey or form, or searching it out themselves from a factual source, like a credit report or public records database.

When asked how companies would figure out characteristics the participant had declined to share with them, participants with this belief felt companies would transparently ask. In response to the interviewer asking how companies would try to learn a user's religion if that user refused to answer a direct question about it:

> "I mean they ask questions and they can somewhat [learn my religion] there. And if they don't, they're gonna ask more questions... If you don't wanna talk about your religion they would probably go…'What type of church do you go to?'...Yeah, ask other questions to try to get around but try to get to the point of whatever it is they're asking about." – MR2

We found that these participants were mostly unaware that companies collect behavioral or other incidental data. Instead, their beliefs hung on largely outdated market research techniques, leaving out automatic or indirect methods of data collection that modern online companies rely on. When we probed whether they believed companies could learn their demographics through a different process, several participants claimed that companies would be unable to learn a characteristic that a user withheld:

> "If you answer that question, it seems like that's what they'll know, that's what they'll have, but if you don't, it seems like they wouldn't know your ethnicity." – MR1

> "I only think that they could figure out my information that I type in." – MR8

Although most focused on companies wanting their demographics, several participants in this cluster also believed that companies are interested in other types of personal information, like addresses, credit card numbers, or social security numbers. They shared stories about personal experiences where their private information was "found out" by companies or individuals searching authoritative sources like public records or credit reports. This method of data collection would be seemingly less visible to the user, but participants who had searched public records for information on themselves or others seemed especially sure that companies would direct their employees to take the same approach. So after obtaining initial information that could seed a search, an employee of the company might look up, for example, a user's age or marital status by seeking out public records.

This is indicative of a common misconception in this cluster about the scale at which companies collect data. It is not feasible for companies to collect data on millions of users by having humans track down public records for each individual, one-by-one. This is, however, the way many in this cluster described companies' data collection processes to us, as humans thumbing through records to find and learn relevant data about an individual:

> "I think they would look at the age. They'd look at the gender. Everything that they have, like where I'm from, where I'm living, what I do, and kind of be like, 'Okay.'" – MR6

This belief that companies make special efforts to directly collect data on each individual was not universal in the MR cluster. Some believed companies simply do not care enough about any single person to hunt down their information, so that companies would ignore and leave alone individuals who decline to share their information. When asked how a company would try to figure out a user's demographics if they withheld it, MR7 said:

> "I think they don't. I think that they just go on. There's so many people. I mean, it's like ants. There's 10,000 of them, and if you kill 9,000, the other 1,000, you're not going to worry about because you got 9,000." – MR7

4.1.2 Companies exploit common sense connections between data to make inferences

The straightforwardness that this cluster ascribed to companies' data collection methods was echoed in their beliefs about how companies analyze data. Market research cluster participants believed companies make inferences by relying not on sophisticated algorithms, but on human employees who make obvious intuitions about the relationship between two pieces of data. The inferences they described companies making were often vague, with their examples tending to revolve around retail recommendations based off of an individual's stated demographics and interests:

> "They ask questions, you answer them, seems like they'll kind of go with whatever you answered. Like if you say you like to ride bikes or something like that, they'll promote bikes, or different things and places you can go to ride bikes. That kind of thing." – MR1

To these participants, companies appear to make an inference based on a single piece of information, and that relationship is intuitive and based on common knowledge. Participants in this cluster did not touch on topics like data aggregation, needing convergent evidence to support an inference, or weak correlations.

A few participants in this cluster did reference retail recommender systems, but their explanations of how these systems work often left out the role of other users' data in guiding recommendations. To some, recommender systems statically present obviously related items as recommendations, e.g., a person buying a phone would be recommended a case for that exact phone. Others believed that companies assume a user's preferences based on their demographics, such as by age or ethnicity.

One misconception about the inferencing process we saw in this cluster was about the directionality of inferences companies make. Although they correctly believed that companies use their characteristics to make inferences about their behavior, many incorrectly believed it was uncommon or impossible for companies to use their behavior to make inferences about their characteristics. When we did prompt them to consider ways that companies might try to infer characteristics from behavior, there was an underlying skepticism that deep insights about a person could come from analyzing online behavior:

> "How could you figure out me by the things I look at?" – MR4

On the contrary, these participants judged companies' inferencing capabilities in terms of their own abilities. We asked participants to explain how companies would infer a characteristic that a user had kept private online, such as their religion or sexual orientation. Participants in this cluster described their own processes as analogues for what they believe companies do, e.g., "While I'm going through somebody's page, I can see a lot about what they're like." (MR8). MR2 put this even more clearly, attaching companies' capabilities to her own:

> "I think they could, 'cause I could." – MR2

As with data collection, these beliefs about inferencing methods appear dated in some respects. Regardless of humans' expertise in making inferences based on intuitive analysis of a person's behavior, companies that serve a large user base have to use inferencing techniques that are scalable in ways that human analysts would not be able to match.

4.1.3 Companies stereotype users based on their demographics, which is morally wrong

The MR cluster included several African-American, Latino, and mixed-ethnicity participants who each expressed concern that inferences companies make appear to be based not on deep knowledge about users but on stereotyping. In their view, companies offer opportunities unequally to people based purely on their ethnicity or income. This was not seen as accidental or benign. Participants who referred to the inferencing process as stereotyping did not mince words. They believed it to be dehumanizing:

> "It begins to be, like, I'm just a statistic for lack of a better word. I'm just a demographic, I'm just a person who spends this amount of money on this in my spare time, and it just becomes - it's so personal but it's impersonal at the same time, you know what I mean? Because it's just information, and they forget that these are people, these are human beings." – MR6

Beyond their moral concerns, they believed stereotyping leads to inaccuracies, particularly due to ignoring intragroup variation:

> "None of us are the same, so we shouldn't be classified as the same...So those companies that put these people in this basket, I think they're sometimes just rounding them up like cattle." – MR4

Those who mentioned this belief were confident that online companies engage in stereotyping, however there was an ambiguity about the exact consequences that result in the examples participants gave of this happening in their own lives:

> "Usually when you do something, they want your background, like your ethnicity or I guess to put you in a certain place, like, maybe they'll know maybe what you want just [based on] your ethnicity. Maybe." – MR1

The ambiguity of the perceived consequences should not obscure the fact that several participants in this cluster believe that this is the process companies engage in. Research on topics like "shopping while Black" [15,34,50] has surfaced how experiences with ambiguous stereotyping are naturally interpreted in light of wider life experiences of racial discrimination, so that online companies' opaque inferencing methods may lead to unflattering interpretations about stereotyping in the absence of clear evidence to the contrary.

4.1.4 High concerns about hackers and scammers can drown out concerns about companies

Although the interviews were only meant to elicit beliefs about companies, several participants in this cluster spontaneously mentioned hackers and scammers as high-stakes threats to their online data. Hackers were described as individuals who would access information either from a device without the owner's knowledge, or via unauthorized access to a company's database. Scammers were described as companies who call, set up phishing websites, or send email in order to obtain information like credit card or social security numbers under false pretenses. The harms participants saw resulting from hackers and scammers were clear: financial loss, identity theft, and damage to their online devices. By comparison, some saw little concrete harm that companies might cause by having their data:

> "You have to worry more about your hackers than you do your companies. Because hackers do bad things with it. They use it, they destroy your credit, they destroy, you know – I don't think a company would want to do that." – MR3

Several participants who shared concerns about companies stereotyping also worried about hackers or scammers misusing their data. These threats appeared to evoke different feelings. Companies stereotyping appeared to create a sense of moral resentment, whereas hackers and scammers came across as adversaries who could be warded off or fought.

4.2 Data Mining Cluster

We now turn to the other main belief cluster. The data mining cluster of beliefs was anchored around a core belief that companies collect data on users' behavior (4.2.1). Participants who had this belief often believed that companies make recommendations based on their prior behavior (4.2.2) (e.g., recommending a song to listen to based on songs the user has previously liked), but they rarely acknowledged that companies can combine demographics with behavioral data to make inferences. They all believed companies use some computer-based processing of user data to make inferences (4.2.3) such as analysis of social connections to make inferences about them (4.2.3), but they varied widely in the role they believed humans play in making inferences. Some believed algorithms work fully independently, whereas others believed that companies have humans-in-the-loop, employees who oversee individual inferences made by algorithms.

Compared to the market research cluster, the data mining cluster was more familiar with implicit data collection methods. Additionally, these participants were more confident in their beliefs about data collection and inferencing, including in their misconceptions. Participants in this cluster often had mixed feelings about data mining (4.2.4), acknowledging the value it may provide to them personally but often exhibiting signs of resignation in the face of little perceived privacy control.

4.2.1 Companies collect users' online behavior data

The participants in this cluster shared the core belief that companies collect data on users' online *behavior*. The exact data mentioned varied by participant but often included links they click, products they purchase, or videos they watch. Unlike the Market Research cluster's belief that companies ask users to purposefully provide data one survey question at a time, these participants felt companies collect implicit behavioral data automatically. They described companies as "collecting", "tracking", or "watching" all of the things they do online. They were aware that companies depend on their behavioral data to provide online services, drawing from experiences when an inferential system explicitly referenced the data it had collected:

> "I go on Amazon a lot, and say I haven't been on in, like, two months. When I log back on, it remembers. It says, 'Oh, you liked this video game,' maybe, 'People who bought this, buy this.'" – DM1

Although they all believed companies collect some kind of behavioral data, they had varied levels of awareness about *how* and *what* behavioral data companies collect. They most commonly mentioned companies saving their history, e.g., searches, purchases, videos watched. Only a few participants mentioned that companies could collect their location, e.g., through GPS, IP address, or searches made. Those who did mention location tracking believed that companies value location data highly due to the variety of inferences they can make from it:

> "My location, for one, is huge. Pretty much everyone wants to use my location...Probably for marketing purposes so that they can use [it] in some way, like your location to establish where you are a lot...What my hobbies are, what stores I go to and shop [at], and basically what I'm doing with my time. Because it could be used for purposes of marketing, I think." – DM3

> "If you go through my location history for, you know, using public transportation, you're gonna know where I work, how I get there, what I do certain days of the week, things like that. I mean, literally, I'm carrying around a tracking device almost 16 hours a day." – DM11

Unlike the Market Research cluster, participants in the Data Mining cluster were generally aware that companies collect incidental browsing data, such as how long they browsed a website, or what type of device they were using to go online. Still, the examples many participants in this cluster gave about companies collecting activity data contained misconceptions. DM7, for instance, knew that companies can aggregate data from across multiple devices if he is signed in to the same account on each one, but he also mistakenly believed that companies can *only* collect data and make inferences about him if he is signed into an account. This could be a costly misconception, as believing that signed out activities cannot be tracked would provide a false sense of privacy online.

4.2.2 Recommendations are based on a user's past behavior

It seemed apparent to participants in this cluster that recommendations of products and online content such as those on retail or social media sites were based on their own past behavior. This was a conclusion that few in the Market Research cluster had come to. The Data Mining participants, on the other hand, shared several examples of recommendations that companies make to them based on their past behavior:

> "I notice a lot of advertisements on my page, especially to sites that I've been to or things that I've looked at." – DM6

> "YouTube makes guesses on me all the time. When I go to YouTube and it shows me things I watched previously, and they'll show [videos they] recommended, so they're always doing that type of stuff." – DM7

DM participants often spoke about repeated interactions with these systems over time providing them insights about how they function. DM2 described her experience with a streaming music service presenting poor recommendations as a result of songs she "liked", leading her to an insight about how the system worked, and how to change her behavior to prevent inaccurate inferences from being made:

> "'Can't Touch This', right? It's that kind of song that [you think], 'Oh, isn't this a cool song?' And you like it. But when they refer songs [based on] that song, it's like, 'Oh, I shouldn't have liked it.' It's like, "Mm, they're going to do something with this, and they're probably going to refer to me stuff [based on] it.' And so I should be wise about what I like." – DM2

Their awareness about a behavioral basis for recommendations does not mean they had fully accurate beliefs about how recommendations are made. One misconception held by some participants in this cluster was that companies would rely only on behavioral data to make recommendations, to the exclusion of other data commonly used in recommendation systems, such as demographics. DM7, for example, believed that companies ignore his age when recommending products or other content:

> "Not too many websites really have shown me things based on my age group." – DM7

4.2.3 Companies use analytics to infer users' characteristics, with varying levels of humans-in-the-loop

Participants in the Data Mining cluster had a common element in their descriptions of how companies make inferences, in that they all had confidence that companies rely on some form of computer-based processing of data:

> "I'm sure there's some kind of algorithm out there, you know, I fall into a certain box, maybe I'm just a number with a letter at the end." – DM11

In that respect, they showed a more accurate perspective on modern inferencing methods than the Market Research cluster, who believed that companies collect the data they are interested in directly. Some participants in this cluster were aware that companies make inferences about them by analyzing their social connections, such as their friends on social media, in addition to their own behavior. However, the Data Mining cluster's other beliefs about inferencing often contained misconceptions about how companies rely on humans or computers to make inferences.

Despite this cluster's belief that computer processing of data is key to inferences, we observed a surprising diversity of beliefs about the role of human oversight in modern inferencing. Some thought that companies rely on automatic processes that make simple connections quickly. These participants talked in terms of computers establishing patterns in a person's behavior:

> "I'm thinking it is a machine scanning somebody's information and kind of learning and getting what they might be interested in or what their habit might be with something." – DM2

> "It makes inferences...I think it's just the computer doing [that], I don't think it's [people]...like keywords, just looking through that, I guess." – DM4

Others believed humans oversee the inferences made by algorithms, micromanaging the process. To these participants, computers are able to generate speculative inferences, but a human would make the final decision about whether an inference is accurate before using it, e.g., to make a recommendation. DM5 believed that humans closely supervise the implementation and results of inferencing programs, potentially leading to inaccuracies based on human judgment:

> "Even though it's electronically collected, electronically manipulated, it's looked at by a human. A human wrote the program. We're fallible." – DM5

Still others in the Data Mining cluster believed that companies rely on employees using computers as shallow tools to aid their own "reasonable skills of deduction", as DM11 put it. DM9 believed companies only use computers to generate visualizations of raw behavioral data, which human analysts would then review to make each inference about each individual user. He felt that companies use this process to determine a person's vulnerabilities, e.g., a person's values or attitudes that can be used to manipulate them via marketing or political appeals:

> "Certainly they've got to have analysts sitting there, you know, they hire interns to get on there and watch all this stuff, 'OK, now put it all down on a chart and show me where is he vulnerable and where is he not.'" – DM9

4.2.4 Mixed or negative feelings about inferencing are common

Participants in this cluster expressed complex feelings about companies making inferences from their behavioral data. This contrasts with some recent work suggesting the privacy calculus that people engage in is more visceral and gut-driven, rather than a purely rational accounting [29]. DM participants often described their use of online services as a trade-off, perceiving both benefits and drawbacks to using online services that rely on their data. DM8, a waste management driver with a keen awareness of behavioral tracking methods, spoke about his decision making process unambiguously, as "does the good outweigh the bad?". Others more broadly considered the purposes that behavioral inferences can serve, contrasting the use of data to save lives against using it for marketing:

> "If we're talking about harming mass quantities of people, like a 9/11 thing, then I'm all for collection of data. But if we're talking about, you know, you want to sell me a crib. {Laughs} Um, then I'm kind of against that." – DM5

Others felt torn as to whether benefits they accrue from inferencing are worth the costs:

> "They would try to tailor something for you specifically for your interest. So I guess one way to look at it is [as an] invasion of privacy and stuff like that. But the other way, you might say, 'Oh that [is] a little bit helpful.' So it's hard to tell." – DM4

Not all participants in this cluster saw advantages to being a part of the data economy. Several participants expressed resignation over their limited ability to control data collection, given that other people can provide data about them online without their consent. DM11, a store clerk in his 30s, was highly concerned about this. He lamented that despite taking strict action to pare down his data footprint by downgrading his smartphone to a feature phone and conscientiously managing his device's privacy settings, he is unable to prevent companies from collecting data about him through his friends' social media posts:

> "The things that I do in real time with real people, they possibly don't have very much access on my end from that. But I can't stop other people from posting things about me on Facebook, Twitter, et cetera." – DM11

We heard several in this cluster speak broadly about data mining as part of what they saw as a general decrease in personal privacy:

> "I'm uncomfortable with it. I didn't sign the deal with the devil basically, aside from hitting yes to a bunch of apps on my shiny, new tablet. Aside from that, I feel it is a great invasion of privacy." – DM11

> "The way everything seems to be going now, it almost seems like there's just less and less privacy...it's just kind of weird, feeling like people know certain things about you, you have no idea...all this information being gathered about you that you don't really know who they are." – DM6

4.3 Comparative analysis between clusters

In addition to the beliefs that defined each cluster, we found several apparent differences in demographics, attitudes, and sense of personal agency between the two clusters. We also include additional information on two misconceptions that spanned both clusters: that companies rely on humans rather than computers to make inferences, and that inferences are made on the basis of a single piece of information.

4.3.1 MR cluster more ethnically diverse; DM had higher income

Although educational attainment was similar across all of our participants, there were important demographic differences between the clusters even in this small sample, including income and ethnicity. Participants with household incomes over $45,000 were almost all in the DM cluster. Each ethnic group in our sample was represented in both clusters, but the MR cluster had a greater proportion of ethnic minority participants compared to the DM cluster: 64% of the DM cluster identified as White or European-American, compared to 38% of the MR cluster. The MR cluster also perceived more stereotyping in companies' behavior, which we return to in the discussion below. Both clusters were roughly equally distributed in terms of age. Although the MR cluster believed that companies engage in older inferencing techniques than the DM cluster, we note that younger participants in the MR cluster had similar beliefs.

4.3.2 DM cluster more specific and confident in their beliefs, including their misconceptions

We observed recurring differences in how participants in each cluster described their beliefs. Compared to the MR cluster, participants in the DM cluster tended to speak more confidently about how they thought companies use their data. MR participants often went out of their way to describe their beliefs as speculative (e.g., "I don't know too much about it, but..." – MR4), but DM participants hedged fewer of their beliefs and misconceptions (e.g., "one way or another, you're being tracked...it happens everywhere" – DM8). Although the MR participants were missing an important piece of the inferencing landscape with regards to behavioral data collection, the DM participants' greater confidence in their misconceptions might be a more difficult obstacle to overcome. We saw hints of this in the teaching intervention section, as the participants with the most confident and specific beliefs during the interview section were often openly resistant to changing their beliefs in response to the learning activities.

4.3.3 MR cluster saw risks, DM cluster saw choices

Both clusters shared what they felt were drawbacks to data collection and inferencing, but they differed in the threats they described and the sense of risk or choice they felt they have in the data economy. We interpret these in light of the demographic differences between the clusters, and how they relate to cultural models of personal agency based on income, and cultural models of interacting with institutions based on ethnicity.

The MR cluster worried whether they are targeted by hackers and scammers, and they felt threatened by what they viewed as companies stereotyping. Many described taking protective measures to guard against what they felt were pervasive threats: financial threats from identity theft and ransomware, or threats to their sense of identity from companies treating them as a stereotype. The language they used often evoked a sense of being under attack, even when the danger was unclear, e.g., "I don't know how it works, but I know I just don't want to be a victim of it." (MR4). This was indicative of what appeared to be a lower sense of personal agency in the data economy in the MR cluster. Even though they felt that the methods companies use to make inferences were not far beyond their own capabilities, we often heard a clear protective motive behind the online behavior this cluster described. The MR participants did not describe trading their information to companies to gain a benefit; instead, they talked protectively about how they tried to prevent their information from being used against them.

The DM cluster had very different concerns and interpretations of their role in decisions about their data. In their view, companies largely provide opportunities for them to consciously trade their data (and by extension, their comfort) for material benefits. Companies appeared in some of their narratives as representing a more abstract threat to the concept of personal privacy, but even those participants felt they are choosing to pay the cost they must, to use the products and services they want:

> "I look at both sides of it and say, 'Well, would I rather do this or would I rather do this?' So if it's not hurting anything, and it could help, then I'm fine with it." – DM8

Somewhat paradoxically, although the MR cluster was more convinced that companies collect their data by explicitly asking them to provide it, the DM cluster seemed to feel more agency and control over the decisions they make online with their data. This difference in perspective may relate to cultural models of agency that differ based on social class, as the DM cluster had higher incomes overall than the MR cluster. We discuss further implications for this finding in the discussion section.

4.3.4 Humans-in-the-loop, and weak correlations in modern inferencing

All of the MR cluster and several in the DM cluster believed that companies make inferences about users by having humans-in-the-loop, either relying on human analysts who exploit common sense connections between two pieces of data (e.g., inferring hobbies from purchases), or by employing experts who analyze an individual's behavior like a detective (e.g., manually combing a user's online pictures for evidence of a spouse). These folk explanations for how companies make inferences exaggerate companies' capabilities in some ways while limiting users' ability to imagine current inferencing methods in others.

Believing that companies rely on common sense logic to tie two data points together ignores the multivariate, deep learning methods that companies now deploy to make unintuitive inferences. To our participants, inferences seemed to be snap judgments based on perfect, intuitive correlations between two pieces of data. This may lead to unpleasant surprises when encountering systems that make unintuitive predictions based on data aggregated from multiple sources. At the same time, the belief that companies employ a team of human experts to diligently analyze each user's data may be partially responsible for some of our users believing that there is no limit to what companies can learn about them. The belief that employees with strong detective skills are hunting down their data may lead some participants to misjudge the risks attached to making different types of data available online for companies to see.

5. DISCUSSION

The patterns we observed in our high school-educated sample's beliefs and misconceptions about companies' inferencing methods underscore the need for privacy researchers to consider qualitative, cultural influences on privacy knowledge, attitudes, and behavior. We share two categories of implications that came out of this work: implications of inference literacy in research, design, and education; and implications of cultural models for future research in online privacy and HCI in general.

5.1 Implications of inference literacy

5.1.1 Redefining tech savviness and digital literacy

As technology itself changes, definitions of tech savviness and digital literacy need to change to stay up-to-date. Measuring tech savviness by the ability to perform instrumental tasks on a local device ignores the extent to which daily device activity takes place in a distributed, online context. Although recent attempts to assess online privacy literacy have gone beyond that to include aspects of how online institutions collect or transmit data [57], digital literacy studies often still rely on self-reported expertise [5,20] or the number of years using the internet [44] as a primary measure of digital literacy. The current study is among the first, to our knowledge, to directly explore this particular aspect of digital literacy: beliefs users have about how companies make inferences from their data.

Our results suggest that inference literacy is worth including as an aspect of digital and online privacy literacy. The overwhelming majority of our participants use multiple devices everyday for various purposes, but they had several misconceptions about current methods of data collection and inferencing that could lead to unpleasant surprises. We advocate for a broadening of the features used to consider what makes a person tech savvy or digitally/privacy literate to include basic inference literacy: (1) that companies can collect and aggregate data from multiple sources including forms, behavioral data, telemetry, and public databases, and (2) that those data are often processed by learning algorithms that can uncover unintuitive or even private connections that can be found due to the massive amount of data available to companies.

5.1.2 The roles of transparency and education in inference literacy

Our participants were active online users who, in the absence of structures to help them build their inference literacy, developed lay theories to explain their online experiences that contained basic misconceptions. We believe this speaks to a need for interventions to support inference literacy, and we discuss potential inference literacy interventions: transparency, as well as informal and formal educational interventions.

First, we consider the issue of transparency. Transparency can inform individuals and surface issues for broader discussion about systemic and policy issues [48]. Users may, for instance, be more comfortable knowing a system does not have humans-in-the-loop when sensitive data are involved, or they may prefer to have humans-in-the-loop if they feel a human could outperform an algorithm. However, transparency is not a silver bullet. It places a heavy burden on the user to learn about the algorithms of each company they engage with, and complex inferential systems are often black boxes even to those who design and deploy them [27]. It may not be feasible to be transparent about inferencing systems that change frequently, and whose true workings require sustained academic research to discover.

Second, there may be a role for informal, "do-it-yourself" interventions that allow users to teach themselves inference literacy concepts [38], such as that aggregating multiple sources of data allows companies to learn unintuitive, weak correlations between data. There are existing resources related to inference literacy that could be used as models for novel interventions. Teachingprivacy.org [69] offers a selection of accessible lessons about online privacy that draw from real world examples, including structured material for deployment by teachers in formal educational settings. Other efforts like R2D3's "A Visual Introduction to Machine Learning" [67] provide more technical knowledge about statistical classification methods. These approaches provide motivated users with resources to correct their misconceptions, but we caution against the idea that these methods will systemically improve inference literacy. Research has shown that relying on "do-it-yourself" approaches to build technology knowledge and skills may widen rather than reduce inequalities in digital literacy [38]. This may be, in part, due to a discoverability issue as a result of jargon used in some informal interventions. Nineteen of our 21 participants had never heard the term "machine learning" prior to the study, which may make it harder for them to find resources like R2D3's.

Finally, we point out that regardless of societal aspirations to increase access to a college education, high school education is still the modal educational attainment in the US. Students graduating with a high school degree should be prepared for more than just college academics; they should also be prepared to live in a world where interactions with inferential systems are common and consequential. Our participants' beliefs were outdated in several crucial respects. The frequent appearance of misconceptions that companies rely on consumers taking surveys to gather demographic data, or on humans-in-the-loop rather than automated analytics to make recommendations, speaks to the danger of assuming that users will osmose the basics of the consumer data ecosystem outside of a formal educational setting. Requiring college or independent study to learn about how personal data may be used to infer a credit score, decide on a loan application, or other important aspects of economic life places that knowledge outside the reach of those who are most likely to be negatively affected by those decisions [14,43]. There is precedent for teaching digital literacy concepts [68] and personal finance [10] at the high school level, and inference literacy is worth considering alongside these topics.

5.2 Implications of cultural models

The current study surfaced several issues related to power and privilege in consumer interactions, which we describe in terms of *cultural models*, sets of assumptions that differ across cultural groups. We share two main insights here: first, that our participants' experiences of risk and choice in online privacy and security relate to cultural models about personal agency that differ by income; second, that our participants from marginalized ethnic groups believed companies' inferencing methods constitute stereotyping, which we link to broader work on ethnic minority experiences in consumer settings. We describe implications of these findings for online privacy research and design, and finish by advocating more broadly for consideration of cultural models as a key lens to critically examine the experiences of marginalized groups in user research.

5.2.1 Income and differences in inference literacy

Educational attainment and income are often treated as equivalent indicators of socioeconomic status, but we saw differences in beliefs and attitudes within our education-controlled sample based on participant income. First, our higher-income participants viewed online privacy decisions as choices they were empowered to make, whereas our lower-income participants framed those same decisions as risks they had to protect against. This echoes prior work showing that middle-class Americans typically develop a sense of agency built around exercising independent choices, whereas working-class Americans often experience greater economic and environmental constraints that preclude such free choice behavior [54,56]. Recently, some inferencing systems have been designed to allow users the ability to modify their workings, either to improve system accuracy or simply to exercise personal choice over their outputs [31,61]. We suspect that users' interactions with these systems may be affected by the larger cultural context in which those choices are made, and we advise system designers to consider how differences in risk perception and personal agency based on cultural differences may affect users' willingness to engage with different system designs.

Second, prior work has found that inequalities in online skills and knowledge often result from differences in SES [20,24,64], and we found a similar, problematic inference literacy gap related to SES. The two different clusters of beliefs we saw were linked to income differences: the MR cluster had nearly all of our participants with under $45,000 household annual income, and was less aware of current data collection and inferencing practices. Although we cannot say whether this trend in our sample is representative of one in the larger population, given that inferential systems already disproportionately negatively affect working-class people [14,43], we again highlight the need for systemic efforts to prevent and reduce digital inequalities, including those related to inference literacy.

5.2.2 Ethnicity and interpretations of inferencing as stereotyping

Several participants spontaneously brought up beliefs that companies stereotype consumers by ethnicity, all of whom claimed that doing so is immoral. It is undoubtedly true that companies use demographics to profile users, and that this is an inherently imprecise process. Modern inferencing systems may include demographics like ethnicity among many features, but these participants believed that inferences are sometimes made based *only* on assumptions about ethnicity. However accurate or inaccurate this belief about stereotyping is, it remains that these participants' life experiences have resulted in a cultural model about interactions with institutions like companies that assumes companies stereotype.

The complex online ecosystem our study explored is often ambiguous as to how decisions are made: the opacity of algorithms that recommend, advertise, or filter content that users see often means users generally lack context for how online companies' decisions and recommendations are made. This leaves plenty of room for the user to interpret online experiences in light of other experiences they have had, including those of being stereotyped or discriminated against. To the user who has experienced discrimination while shopping [15,34,50], driving [37], or merely walking [16], stereotyping by online companies may appear no different. Designers should therefore take caution in how they include or describe ethnicity as a component of decision-making about users. Lacking clear evidence to the contrary, unflattering interpretations may be made about inferential systems for which the purpose of using demographics like ethnicity is left ambiguous to the user.

5.2.3 Cultural models in user research and design

In this work, we applied the concept of cultural models to describe additional layers of commonalities and differences across our participants' experiences. Although the finding that our two clusters had different views on risk and choice online might stand on its own, incorporating cultural models allowed us to link this finding to different beliefs about personal agency that relate to social class rather than leaving our analysis at the level of the individual participant. This provided us insight into a mechanism that seems to underlie interpretations about online privacy consequences, one that speaks to different economic and environmental constraints between cultural groups. We believe that exploring the ways that cultural models qualitatively affect people's interpretations and attitudes about online phenomena complements other user research approaches by providing a textured, layered perspective on the meaning that users attach to their online privacy experiences.

5.3 Future Directions

We advocate for further research on inference literacy in high school- and college-educated samples to confirm whether the belief clusters we observed exist in the larger population, as well as to further explore whether inference literacy varies by educational attainment or geographic location. We also endorse the adoption of cultural models as a useful lens to apply to other research in online privacy. It would also be valuable to further explore how inference literacy beliefs interact with participants' online behavior and decision-making processes, in order to inform new system designs that can better support inference literacy.

5.4 Conclusion

We began this work by describing the vast difference between companies' past and present inferencing methods. There is little reason to believe that current methods will remain static, but our findings suggest that there is already a substantial gap between what people believe companies are doing with their data, and the current reality of pervasive, automatic algorithms. We point not only to the size of that gap, but also to its heterogeneity: we saw links between inference literacy beliefs and larger cultural models based on income, ethnicity, and potentially educational attainment. Culturally sensitive policy, research, and design may be a route to minimizing digital inequalities that arise as an outcome of group differences in inference literacy.

6. ACKNOWLEDGMENTS

We thank our participants for their contributions to this research. We also thank Sunny Consolvo and Tara Matthews for their valuable guidance. We express our gratitude to Paul Aoki, Erik Ninomiya, Katie O'Leary, Sai Teja Peddinti, Shepherd Pittman, Sonam Samat, Martin Shelton, and Irene Tang for their helpful feedback and assistance.

7. REFERENCES

[1] Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). "Privacy and human behavior in the age of information," *Science, 347*(6221), 509–514.

[2] Alpaydin, E. (2014). *Introduction to Machine Learning,* MIT Press.

[3] Aronson, J., & Inzlicht, M. (2004). "The ups and downs of attributional ambiguity stereotype vulnerability and the academic self-knowledge of African American college students," *Psychological Science, 15*(12), 829–836.

[4] Awad, M., & Khanna, R. (2015). "Machine learning and knowledge discovery," in *Efficient Learning Machines: Theories, Concepts, and Applications for Engineers and System Designers*, Berkeley, CA: Apress.

[5] Bartsch, M., & Dienlin, T. (2016). "Control your Facebook: An analysis of online privacy literacy," *Computers in Human Behavior, 56*, 147–154.

[6] Berkovsky, S., & Freyne, J. (2010). "Group-based recipe recommendations: analysis of data aggregation strategies," *In Proceedings of the 4th ACM Conference on Recommender Systems: RecSys '10*, 111–118.

[7] Bi, B., Shokouhi, M., Kosinski, M., & Graepel, T. (2013). "Inferring the demographics of search users: Social data meets search queries," *In Proceedings of the 22nd International Conference on World Wide Web: WWW '13*, 131–140.

[8] Braun, V., & Clarke, V. (2006). "Using thematic analysis in psychology," *Qualitative Research in Psychology, 3*(2), 77–101.

[9] Choi, B., Lee, I., Kim, J., & Jeon, Y. (2005). "A qualitative cross-national study of cultural influences on mobile data service design," *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: CHI '05*, 661–670.

[10] Council for Economic Education. (2016). "Survey of the states: Economic and personal finance education in our nation's schools", Retrieved from Council for Economic Education: http://councilforeconed.org/wp/wp-content/uploads/2016/02/sos-16-final.pdf

[11] D'Andrade, R. (1995). *The Development of Cognitive Anthropology,* New York, NY: Cambridge University Press.

[12] van Deursen, A. J. A. M., van Dijk, J. A. G. M., & ten Klooster, P. M. (2015). "Increasing inequalities in what we do online: A longitudinal cross sectional analysis of Internet activities among the Dutch population (2010 to 2013) over gender, age, education, and income," *Telematics and Informatics, 32*(2), 259–272.

[13] Dinev, T., Masssimo, B., Hart, P., Christian, C., Vincenzo, R., & Ilaria, S. (2005). "Internet users, privacy concerns and attitudes towards government surveillance: An exploratory study of cross-cultural differences between Italy and the United States," *In Proceedings of Bled: BLED '05*.

[14] Fourcade, M., & Healy, K. (2013). "Classification situations: Life-chances in the neoliberal era," *Accounting, Organizations and Society, 38*(8), 559–572.

[15] Gabbidon, S. L. (2003). "Racial profiling by store clerks and security personnel in retail establishments: An exploration of 'shopping while Black,'" *Journal of Contemporary Criminal Justice, 19*(3), 345–364.

[16] Gelman, A., Fagan, J., & Kiss, A. (2007). "An analysis of the New York City Police Department's 'stop-and-frisk' policy in the context of claims of racial bias," *Journal of the American Statistical Association, 102*(479), 813–823.

[17] Golbeck, J., Robles, C., Edmondson, M., & Turner, K. (2011). "Predicting personality from Twitter," *In Proceedings of Privacy, Security, Risk and Trust and International Conference on Social Computing: PASSAT/SocialCom '11,* 149–156.

[18] Gou, L., Zhou, M. X., & Yang, H. (2014). "KnowMe and ShareMe: understanding automatically discovered personality traits from social media and user sharing preferences," *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: CHI '14*, 955–964.

[19] Harboe, G., & Huang, E. M. (2015). "Real-world affinity diagramming practices: Bridging the paper-digital gap," *In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems: CHI '15,* 95–104.

[20] Hargittai, E., & Hinnant, A. (2008). "Digital inequality: Differences in young adults' use of the internet," *Communication Research, 35*(5), 602–621.

[21] Heimgärtner, R. (2013). "Reflections on a model of culturally influenced human–computer interaction to cover cultural contexts in HCI design," *International Journal of Human-Computer Interaction, 29*(4), 205–219.

[22] Helgeson, J. G., Kluge, E. A., Mager, J., & Taylor, C. (1984). "Trends in consumer behavior literature: A content analysis," *Journal of Consumer Research*, *10*(4), 449–454.

[23] Holland, D., & Quinn, N., eds. (1987). *Cultural Models in Language and Thought,* New York, NY: Cambridge University Press.

[24] Hsieh, J. J. P.-A., Rai, A., & Keil, M. (2008). "Understanding digital inequality: Comparing continued use behavioral models of the socio-economically advantaged and disadvantaged," *MIS Quarterly, 32*(1), 97–126.

[25] Ion, I., Reeder, R., & Consolvo, S. (2015). "'... no one can hack my mind': Comparing Expert and Non-Expert Security Practices," *In Proceedings of Symposium On Usable Privacy and Security: SOUPS '15*, 327–346.

[26] Kang, R., Dabbish, L., Fruchter, N., & Kiesler, S. (2015). "'My data just goes everywhere:' User mental models of the internet and implications for privacy and security," *In Proceedings of Symposium On Usable Privacy and Security: SOUPS '15,* 39–52.

[27] Kapoor, A., Lee, B., Tan, D., & Horvitz, E. (2010). "Interactive optimization for steering machine classification," *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: CHI '10*, 1343–1352.

[28] Kauer, M., Günther, S., Storck, D., & Volkamer, M. (2013). "A comparison of American and German folk models of home computer security," *In Proceedings of Human Aspects of Information Security, Privacy, and Trust: HAS '13*, 100–109.

[29] Kehr, F., Kowatsch, T., Wentzel, D., & Fleisch, E. (2015). "Blissfully ignorant: The effects of general privacy concerns, general institutional trust, and affect in the privacy calculus" *Information Systems Journal, 25*(6), 607–635.

[30] Kosinski, M., Stillwell, D., & Graepel, T. (2013). "Private traits and attributes are predictable from digital records of human behavior," *PNAS, 110*(15), 5802–5805.

[31] Kulesza, T., Stumpf, S., Burnett, M., & Kwan, I. (2012). "Tell me more?: The effects of mental model soundness on personalizing an intelligent agent," *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: CHI '12*, 1–10.

[32] Kulesza, T., Stumpf, S., Burnett, M., Yang, S., Kwan, I., & Wong, W.-K. (2013). "Too much, too little, or just right? Ways explanations impact end users' mental models," *In Proceedings of IEEE Symposium on Visual Languages and Human Centric Computing: VL/HCC'13*, 3–10.

[33] Kurasaki, K. S. (2000). "Intercoder reliability for validating conclusions drawn from open-ended interview data," *Field methods, 12*(3), 179–194.

[34] Lee, J. (2000). "The salience of race in everyday life: Black customers' shopping experiences in Black and White neighborhoods," *Work and Occupations, 27*(3), 353–376.

[35] Lewenberg, Y., Bachrach, Y., & Volkova, S. (2015). "Using emotions to predict user interest areas in online social networks," *In Proceedings of IEE International Conference on Data Science and Advanced Analytics: IEEE DSAA '15*.

[36] Li, Y. (2012). "Theories in online information privacy research: A critical review and an integrated framework," *Decision Support Systems, 54*(1), 471–481.

[37] Lundman, R. J., & Kaufman, R. L. (2003). "Driving while Black: Effects of race, ethnicity, and gender on citizen self-reports of traffic stops and police actions," *Criminology, 41*(1), 195–220.

[38] Matzat, U., & Sadowski, B. (2012). "Does the 'do-it-yourself approach' reduce digital inequality? Evidence of self-learning of digital skills," *The Information Society, 28*(1), 1–12.

[39] McDonald, A. M., & Cranor, L. F. (2010). "Americans' attitudes about internet behavioral advertising practices," *In Proceedings of the 9th Annual ACM Workshop on Privacy in the Electronic Society: WPES '10*, 63–72.

[40] Miller, P. J., Cho, G. E., & Bracey, J. R. (2005). "Working-class children's experience through the prism of personal storytelling," *Human Development, 48*(3), 115–135.

[41] Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective,* Cambridge, MA: MIT Press.

[42] Oakleaf, M. (2009). "The information literacy instruction assessment cycle: A guide for increasing student learning and improving librarian instructional skills," *Journal of Documentation, 65*(4), 539–560.

[43] Pager, D., & Shepherd, H. (2008). "The sociology of discrimination: Racial discrimination in employment, housing, credit, and consumer markets," *Annual Review of Sociology, 34*, 181.

[44] Park, Y. J. (2011). "Digital literacy and privacy behavior online," *Communication Research*, *40*(2), 215-236.

[45] Park, Y. J. (2013). "Offline status, online status: Reproduction of social categories in personal information skill and knowledge," *Social Science Computer Review, 31*(6), 680–702.

[46] Pearce, K. E., & Rice, R. E. (2013). "Digital divides from access to activities: Comparing mobile and personal computer internet users", *Journal of Communication, 63*(4), 721–744.

[47] Pearce, R. R. (2006). "Effects of cultural and social structural factors on the achievement of White and Chinese American students at school transition points," *American Educational Research Journal, 43*(1), 75–101.

[48] Rader, E. (2014). "Awareness of behavioral tracking and information privacy concern in Facebook and Google," *In Proceedings of Symposium On Usable Privacy and Security: SOUPS '14*.

[49] Rader, E., Wash, R., & Brooks, B. (2012). "Stories as informal lessons about security," *In Proceedings of Symposium On Usable Privacy and Security: SOUPS '12*.

[50] Schreer, G. E., Smith, S., & Thomas, K. (2009). "'Shopping while Black': Examining racial discrimination in a retail setting," *Journal of Applied Social Psychology, 39*(6), 1432–1444.

[51] Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Ramones, S. M., Agrawal, M., Shah, A., Kosinski, M., Stillwell, D., Seligman, M. E. P., & Ungar, L. H. (2013). "Personality, gender, and age in the language of social media: The open-vocabulary approach," *PLoS ONE, 8*(9).

[52] Sellers, R. M., & Shelton, J. N. (2003). "The role of racial identity in perceived racial discrimination," *Journal of Personality and Social Psychology, 84*(5), 1079–1092.

[53] Smith, A. (2015). "U.S. smartphone use in 2015," Pew Research Center: Internet, Science & Tech. Retrieved from Pew Research Center: http://www.pewinternet.org/2015/04/01/us-smartphone-use-in-2015/

[54] Snibbe, A. C., & Markus, H. R. (2005). "You can't always get what you want: Educational attainment, agency, and choice," *Journal of Personality and Social Psychology, 88*(4), 703–720.

[55] Steele, C. M., & Aronson, J. (1995). "Stereotype threat and the intellectual test performance of African Americans," *Journal of Personality and Social Psychology, 69*(5), 797.

[56] Stephens, N. M., Fryberg, S. A., & Markus, H. R. (2012). "It's your choice: How the middle-class model of independence disadvantages working-class Americans," *Facing Social Class: How Societal Rank Influences Interaction*, New York, NY: Russell Sage Foundation.

[57] Trepte, S., Teutsch, D., Masur, P. K., Eicher, C., Fischer, M., Hennhöfer, A., & Lind, F. (2015). "Do people know about privacy and data protection strategies? Towards the 'Online Privacy Literacy Scale'(OPLIS)," *Reforming European Data Protection Law*, Dordrecht, NL: Springer, 333–365.

[58] Ur, B., Leon, P. G., Cranor, L. F., Shay, R., & Wang, Y. (2012). "Smart, useful, scary, creepy: Perceptions of online behavioral advertising," *In Proceedings of Symposium On Usable Privacy and Security: SOUPS '12*.

[59] Ur, B., & Wang, Y. (2013). "A cross-cultural framework for protecting user privacy in online social media," *In Proceedings of the 22nd International Conference on World Wide Web: WWW '13*, 755–762.

[60] U.S. Census Bureau. (2015). Educational Attainment in the United States: 2014 - Detailed Tables. Retrieved from U.S. Census Bureau: http://www.census.gov/hhes/socdemo/education/data/cps/2014/tables.html

[61] Warshaw, J., Matthews, T., Whittaker, S., Kau, C., Bengualid, M., & Smith, B. A. (2015). "Can an algorithm know the 'real you'?: Understanding people's reactions to hyper-personal analytics systems," *In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems: CHI '15*.

[62] Wash, R. (2010). "Folk models of home computer security," *In Proceedings of the Symposium on Usable Privacy and Security: SOUPS '10*.

[63] Wash, R., & Rader, E. (2015). "Too much knowledge? security beliefs and protective behaviors among United States internet users," *In Proceedings of the Symposium on Usable Privacy and Security: SOUPS '15*.

[64] Wei, L. (2012). "Number matters: The multimodality of internet use as an indicator of the digital inequalities," *Journal of Computer-Mediated Communication, 17*(3), 303–318.

[65] Wei, L., & Hindman, D. B. (2011). "Does the digital divide matter more? Comparing the effects of new media and old media use on the education-based knowledge gap," *Mass Communication and Society, 14*(2), 216–235.

[66] Weinsberg, U., Bhagat, S., Ioannidis, S., & Taft, N. (2012). "BlurMe: Inferring and obfuscating user gender based on ratings," *In Proceedings of the 6th ACM Conference on Recommender Systems: RecSys '12*, 195–202.

[67] Yee, S., & Chu, T. (2015). "A Visual Introduction to Machine Learning". Retrieved from R2D3: http://www.r2d3.us/visual-intro-to-machine-learning-part-1/.

[68] "DigitalLiteracy.gov: Your destination for digital literacy resources and collaboration," Retrieved from National Telecommunications and Information Administration: http://www.digitalliteracy.gov/

[69] "Teaching Privacy". Retrieved from International Computer Science Institute: http://teachingprivacy.org/

Appendix

Appendix A. Teaching intervention procedure

The teaching intervention included a pre-test assessment; two interventions to teach inference literacy concepts, each followed by a card sorting task to assess users' developing explanations of inferencing phenomena; and a post-test assessment to gauge changes in inferencing beliefs after the interventions.

The **pre-test assessment** consisted of one main prompt and follow-up questions about the inferencing capabilities of a social media company.

Next, we provided the **first teaching intervention**, in which the interviewer explained that companies may collect behavioral data while people use a device. Following this first intervention, the interviewer gave the participant the **first card sorting task**, in which participants ranked the likelihood that a given inference could be drawn from a particular piece of data, e.g., "Data: List of apps on phone, Inference: Whether they have kids". These inferences were chosen because they could be made intuitively by people or by an algorithm, allowing us to learn which explanation participants gravitated towards. Afterwards, the interviewer provided feedback on which inferences are or are not likely to be possible for companies to make using current technology.

Next, we provided the **second teaching intervention**, sharing a simplified explanation of classification through machine learning. After the second teaching intervention, we provided the **second card sorting task**, with inferences chosen to explore capabilities related to classification, e.g., "Data: Text from social media posts, Inference: Their personality" [14]. We hoped that after the explanation of classification, participants' explanations would include details of the machine learning process. Again, the interviewer provided feedback on the feasibility of each inference.

Finally, the interviewer administered a **post-test assessment**, asking about the inferencing capabilities of a cell phone service provider. We finished the session by soliciting feedback on the teaching intervention, which we used to refine materials between cycles.