A User-Centered Perspective on Algorithmic Personalization

Rena Coen, UC Berkeley
Emily Paul, UC Berkeley
Pavel Vanegas, UC Berkeley
Alethea Lange, Center for Democracy & Technology
G.S. Hans, Center for Democracy & Technology

May 6, 2016

Master of Information Management and Systems: Final Project

School of Information

University of California, Berkeley

Berkeley, CA 94720

A User-Centered Perspective on Algorithmic Personalization

Rena Coen, UC Berkeley School of Information; Emily Paul, UC Berkeley School of Information; Pavel Vanegas, UC Berkeley School of Information

with Alethea Lange, Center for Democracy & Technology; G.S. Hans, Center for Democracy & Technology

I. INTRODUCTION …… 3
II. BACKGROUND …… 5
- A. Data Practices Approach to Personalization …… 6
- B. Values-Based Approach to Personalization …… 8
- C. Empirical Work on Personalization …… 10
- D. Research Questions …… 12

III. METHODS …… 13
- A. Instrument …… 13
- B. Survey Participants …… 14
- C. Dependent Variables …… 16
- D. Independent Variables …… 16

IV. RESULTS …… 17
- Limitations …… 21

V. DISCUSSION OF FINDINGS …… 22
- A. Personal Information …… 22
- B. Contextualizing User Attitudes About Personalization …… 23
  - Data Type Sensitivity …… 23
  - Avoid Personalization Based on Race and Income Level …… 23
  - Gender Should Be Used With Care …… 24
  - City or Town of Residence Is Relevant for Some Personalization …… 25

VI. CONCLUSION …… 27
VII. ACKNOWLEDGMENTS …… 29
VIII. APPENDIX A - VIGNETTES …… 30
IX. APPENDIX B - RESULT TABLES …… 33

A User-Centered Perspective on Algorithmic Personalization

Rena Coen, UC Berkeley School of Information; Emily Paul, UC Berkeley School of Information; Pavel Vanegas, UC Berkeley School of Information

with Alethea Lange, Center for Democracy & Technology; G.S. Hans, Center for Democracy & Technology

INTRODUCTION

This study focuses on user attitudes toward algorithmic personalization, or the process by which user data is employed to tailor content to users online. Much of the content we encounter online is individually tailored. The Internet is unique to every user, from advertising to search results and even sometimes the prices we are offered. This personalization can save time by extrapolating unstated details, like our location, to enrich the quality of our online experience. For example, Google shows us results from the local theater when we search "Star Wars showtime." This is part of the magic of the Internet, but the technology is not always welcome and may not benefit all users equally. Some forms of personalization may be seen as unfair because it excludes them from content or experiences online or draws on sensitive personal data. Personalization can also disadvantage some groups when it is predicated on negative assumptions or biased towards the preferences of the statistical majority. The concerns raised by those critical of big data analytics are amplified by the relative obscurity of personalization technology. Personalization is typically hidden from users, limiting their ability to object to (or express any opinion on) how they are steered around the Web. Choices are made behind the scenes—predicting what will be of interest to a particular user—that influence what content is prioritized or shown at all.

The technology powering personalization is based on on an ever-broadening array of personal data gathered and aggregated from different sources. The computational capacities that facilitate the collection, storage, and analysis of these large datasets—often referred to simply as “big data”—have facilitated this growth in automated, or algorithmic personalization. In its most basic form, an algorithm is a set of step-by-step instructions—a recipe—“that leads its user to a particular answer or output based on the information at hand.”[1] In the context of automated decision-making, an algorithm can calculate a prediction, a characterization, or an inferred attribute, all of which can then be used as the basis for personalization. This basic concept can be deployed with varying degrees of sophistication, powered by big data’s large volume and diverse variety of data, as well as the rapid velocity with which it moves—the “3 Vs.”[2]

[1] Christopher Steiner, “Automate this: How algorithms came to rule our world,” (New York: Portfolio/Penguin, 2012), 5.

[2] Executive Office of the President, “Big Data: Seizing Opportunities, Preserving Values,” White House, 2014, 4.

Many scholars and advocates, as well as regulators and policymakers, have raised concerns about the discriminatory potential of decisions supported by big data.[3] Much of this work focuses on the use of big data in decision-making contexts that are covered by anti-discrimination law, such as housing, employment, and credit, and looks to these existing anti-discrimination laws as a starting point to address big data's discriminatory potential. In addition to these more regulated domains, automated decision-making affects people in newer, less regulated contexts: through many of the seamlessly automated exchanges of information that take place as people use the Internet.

In this study we do not dig deeply into the technical mechanisms used to support algorithmic personalization. Rather, our focus is on how users feel about the effects of online personalization, we look in particular at the use of consumer data to tailor the content that is shown to them. Drawing from a large-scale survey which includes experimental vignettes, we aim to bring a user-centered perspective to existing policy discussions on online personalization practices and how the potential harms from these practices should be addressed. We observe how users feel about the use of inferred personal information to support personalization, and whether the accuracy of those inferences is important to them. We take a contextual approach to understanding user attitudes about these aspects of online personalization by looking at them in three different domains: advertisements, search results, and pricing, all with a range of types of personal data.

We find some evidence that users value accuracy in the inferences used to support personalization. However, we find that user attitudes about personalization are highly dependent on the context in which it takes place (in our study, targeted ads, filtered search results, or differential pricing), as well as on the basis for that personalization (e.g. location, gender, or race). The accuracy of inferences matters more when the personalization is viewed as neutral or fair while in cases where users see the use of the data type to personalize in a given context as unfair, the accuracy of the inference does not improve user attitudes. We suggest that users have positive attitudes about personalizations that reflect familiar practices, such as prices that vary based on city or town of residence. Finally, we find that the use of race and household income level as the basis for personalization is viewed negatively by users across all contexts.

These empirical results point to the fact that decisions about the implementation of personalizations cannot be made solely based on a general perception of the sensitivity of the data type. User attitudes may depend on the domain in which the personalization occurs and its effects as well as on the perceived relevance of the data to the personalization. Our findings

---

[3] E.g., Solon Barocas and Andrew D. Selbst, "Big Data's Disparate Impact," *California Law Review*, Vol. 104, 2016. Available at SSRN: http://ssrn.com/abstract=2477899; Federal Trade Commission, "Big Data: A Tool for Inclusion or Exclusion," 2016; Oscar Gandy, "Engaging Rational Discrimination: Exploring Reasons for Placing Regulatory Constraints on Decision Support Systems," *Ethics Inf. Technol.* (2010) 12:29-42.

suggest that some personal data types (whether inferred or provided) should not be used as the basis for personalization at all. In cases where personalization may be seen as fair and useful to consumers, companies should implement data practices that are conscious not only of the process-based methods of the Fair Information Practice Principles (FIPPS), but also of the ways that the context of the personalization may violate cultural values, resulting in potentially negative effects on user attitudes.

BACKGROUND

Online personalization is the use of information about an individual user or group of users to tailor website content to them. We define algorithmic online personalization broadly, acknowledging that the source of the information that serves as the basis for the personalization and the logic of the tailoring can be automated to varying degrees. Using information that has been collected or inferred as a basis, the decision about what content to target users with can be based on more traditional deductive analysis, be fully automated through machine learning, or be a hybrid of the two.

The FTC report *Big Data: A Tool for Inclusion or Exclusion?* outlines four phases in the life cycle of big data: "(1) collection; (2) compilation and consolidation; (3) analysis; and (4) use." [4,5] Data is first collected from consumers by tracking their online and offline activity as well as by relying on consumers to provide it directly when they log into services or purchase items. Compilation and consolidation involves assembling profiles about consumers, which is often done by data brokers. The analysis phase uses statistical models to identify patterns and make inferences about consumers including their interests and likely future purchases. In the fourth phase, these analyses are then used to market to consumers or otherwise tailor content to them. Companies have long analyzed consumer data to inform their business practices, but big data presents increasing opportunities to draw on a wide variety and large volume of consumer data. Companies are also now able to fully automate this analysis and use, with increasingly complex and opaque outcomes for users.

In some cases, companies may intentionally infer demographic or other personally identifiable information about users. For example, Target's identification of pregnant women for ad targeting [6] and Universal Pictures and Facebook's targeted ad campaign for *Straight Outta Compton* based on a user's inferred ethnic affinity group.[7] In other cases, companies may not explicitly

[4] Federal Trade Commission, "Big Data: A Tool for Inclusion or Exclusion," 2016, 3.
[5] See: Daniel J. Solove, "A Taxonomy of Privacy," *University of Pennsylvania Law Review*, Vol. 154, No. 3, p. 477, January 2006; GWU Law School Public Law Research Paper No. 129. Available at SSRN:http://ssrn.com/abstract=667622 (Each of these phases can give rise to privacy harms. There are some parallels to and overlap with Solove's taxonomy of privacy in which he categorizes harmful activities in four groups: 1) information collection; 2) information processing; 3) information dissemination; and 4) invasion.)
[6] Charles Duhigg, "How Companies Learn Your Secrets," *New York Times*, February 16, 2012.
[7] Nathan McAlone, "Here's why 'Straight Outta Compton' had different Facebook trailers for people of different races," *Business Insider*, March 16, 2016.

personalize based on demographic information but instead use activity or location-based data that may act as a proxy for these sensitive traits. In 2012, a *Wall Street Journal* investigation found that Staples.com was showing people different prices for the same products. Staples did not disclose the factors used to tailor prices, but the investigation found the strongest correlation between price and the distance to a competitor store from the users' ZIP codes.[8] While we don't know the intentions of the analysts who wrote this program, in this case, higher prices were being shown to users in lower income areas. Whether or not socioeconomic status was factored into the equation directly (fundamentally a data process question), is insufficient to address the ultimately discriminatory outcome. ZIP code often correlates with factors that are considered unfair or discriminatory such as race and income due to historical patterns of discrimination. The negative response to the effects of the differential pricing on Staples.com demonstrates that the outcomes matter and that personalization may be seen as unfair and discriminatory even if it is not intentionally targeting a particular group.

Personalization raises questions about fairness online, both based on data practices and with respect to the norms and values that underpin privacy frameworks. While the data practice and privacy literature (and corresponding best practices of each) are not entirely separate, they have distinct perspectives. Some have focused on the procedures of data collection and use as a place to intervene on behalf of individual interests and integrate fairness into the process. However, others argue for a values-driven solution that responds to the perception the user has of the practice and its results. The difference in these approaches is not necessarily in their effect, but in their underlying philosophy. We look to both the process-based approach to data practices and to the values-based approach to privacy to interpret user attitudes about personalization, and consider existing policies and proposed solutions through a user-centered lens. The process-based data practices approach allows us to scrutinize the mechanisms by which companies obtain consumer data, how they treat it, and the decisions they implement based on that data, while the privacy literature allows us to bring in the values-based discussion of norms and how users feel when certain types of personal data are obtained by different actors and used in different contexts.

Data Practices Approach to Personalization

The Federal Trade Commission is central to discussions of fairness in data practices. In its 2012 report *Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers,*[9] the FTC outlined a framework of best practices for the collection and use of consumer data. This framework is based on the Fair Information Practice Principles (FIPPs) and is designed to "be useful to companies as they develop and maintain processes and systems to operationalize privacy and data security practices within their businesses."[10] The

[8] Jennifer Valentino-Devries, Jeremy Singer-Vine, and Ashkan Soltani, "Websites Vary Prices, Deals Based on Users' Information," *Wall Street Journal*, December 24, 2012.
[9] Federal Trade Commission, "Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers" (2012).
[10] Ibid., 1

framework is intended for use by businesses that "collect or use consumer data that can be reasonably linked to a specific consumer, computer, or other device, unless the entity collects only non-sensitive data from fewer than 5,000 consumers per year and does not share the data with third parties."[11] This scope extends beyond data traditionally seen as personally identifiable by including data that can be linked to a device. This addresses a key issue in online personalization in which a website does not need to identify an individual in order to tailor content to them. The FTC acknowledges that potential privacy harms extend beyond economic or physical harm and unwarranted intrusions to include the "unexpected revelation of previously private information, including both sensitive . . . and less sensitive information to unauthorized third parties."[12]

The FTC report notes that many companies already use the FIPPs to inform their business practices. Like the FIPPs, the FTC framework emphasizes accuracy in its principle that "companies should incorporate substantive privacy protections into their practices, such as data security, reasonable collection limits, sound retention and disposal practices, and data accuracy."[13] The Commission calls for a flexible approach to accuracy that is "scaled to the intended use and sensitivity of the information."[14] They differentiate between data that is used for marketing purposes which they say does not require the same level of accuracy as data used to determine "consumers' eligibility for benefits."[15] The Data Quality & Integrity principle of the FIPPs similarly includes accuracy as a key component of the fair handling of consumer data. While the inference of user characteristics raises concerns about accuracy, it also allows companies to base personalization on data without user awareness or consent.

Many of the proposals to address the harms of inference emphasize the importance of accuracy and of allowing users to correct inaccurate data about them. Work by Crawford and Schultz,[16] Citron and Pasquale,[17] and others proposes due process as a means of addressing and mitigating the harms of automated decision-making. Crawford and Schultz recommend that "those who use Big Data to 'adjudicate' others [be required] to post some form of notice, disclosing not only the type of predictions they attempt, but also the general sources of data that they draw upon as inputs."[18]

---

[11] Ibid., 15
[12] Ibid., 8
[13] Ibid., 30
[14] Ibid., 30
[15] Ibid., 30
[16] Kate Crawford and Jason Schultz, "Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms," *Boston College Law Review*, Vol. 55, No. 93, (2014), Available at SSRN: http://ssrn.com/abstract=2325784.
[17] Danielle Keats Citron and Frank Pasquale, "The Scored Society: Due Process for Automated Predictions." *89 Washington Law Review 1* (2014).
[18]Kate Crawford and Jason Schultz, "Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms," *Boston College Law Review*, Vol. 55, No. 93, (2014), Available at SSRN: http://ssrn.com/abstract=2325784, 125.

The FIPPs, as well as the White House and FTC frameworks that draw on them, emphasize users' involvement in the collection and use of their data as a procedural mechanism to ensure fairness. However, as discussed in the privacy literature cited below, inference raises direct challenges to these existing mechanisms because it involves the creation of unexpected new data about consumers. Arguably, the FIPPs does not explicitly cover or contemplate data generated or inferred from other data. As a result, given how companies may adopt personalization by using inference, the FIPPs may be an insufficient mechanism for protecting privacy.

Understanding how users respond to the use of inference and how aware they are of it can help guide the implementation of the data practice frameworks designed to ensure consumers are treated fairly. The need to understand consumer expectations and attitudes is apparent in the FTC's principle stating that "companies do not need to provide choice before collecting and using consumer data for practices that are consistent with the context of the transaction or the company's relationship with the consumer."[19] In its 2012 Consumer Privacy Bill of Rights, the White House includes a similar Respect for Context principle and acknowledges that "research on consumers' attitudes and understandings"[20] is necessary to apply this principle in a context-sensitive manner.

Values-Based Approach to Personalization

Privacy scholars have raised the issues that we look at in this research around personalization and inference through a values and norms-based lens that focuses on the context and effects of data use. The importance of contextual understandings of privacy has been discussed extensively, including by Helen Nissenbaum, who proposed the concept of contextual privacy norms.[21] This contextual approach to privacy has been tested in empirical research by Shilton and Martin, who found that users' privacy expectations for mobile applications varied based on data type and social context.[22]

Privacy harms can potentially arise from the collection or inference of personal data as well as from the use of that data to personalize content. In their book chapter "Big Data's End Run around Anonymity and Consent," Barocas and Nissenbaum argue that in attempting to address big data's threats to privacy, "procedural approaches cannot replace policies based on

---

[19] Federal Trade Commission, "Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers" (2012)., 48

[20] White House. "Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy," (2012), 16.

[21] Helen Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. (Redwood City, CA: Stanford University Press, 2009).

[22] Katie Shilton and Kirsten E. Martin, "Mobile Privacy Expectations in Context," TPRC 41: The 41st Research Conference on Communication, Information and Internet Policy. (March 24, 2013), Available at SSRN:http://ssrn.com/abstract=2238707 or http://dx.doi.org/10.2139/ssrn.2238707.

substantive moral and political principles that serve specific contextual goals and values."[23] They insist that privacy is central to protection against the harms of big data but that in the context of big data, privacy can not be conceived of as control or secrecy. Using the conception of privacy as appropriate flows of information, their evaluation of the fairness of big data practices is based on whether the practices conflict with information-flow norms and expectations, thereby violating contextual integrity. They advocate evaluating big data practices based on whether they promote context-relevant values.

Dwork and Mulligan's 2013 paper "It's Not Privacy and It's Not Fair,"[24] argues that the focus on privacy and transparency in the discussions of automated decision-making and in proposed solutions to the negative effects of classification fails to address the key issues. They posit that users may experience classifications and the decisions made based on them as unfair even if the data used to support them was obtained with users' permission. This work draws an important distinction between the harms arising from the inferences that often support personalization and the personalization itself.

Algorithmic decision-making involves inference, which can occur at the analysis phase, the use phase, or, as is often the case, both. Inference may be used in the analysis phase to classify people into different categories; these can include things we traditionally think of as personal information, such as gender, education level, and income, as well as marketing segments or interest-based categories.[25] In other cases, directly collected (i.e. not inferred) data may be used in pattern-detection algorithms, resulting in a targeting decision based on similarities between the current user and, for example, people who have bought a product in the past. Often, these two steps are conflated and there is not an explicit inference or decision. For the purposes of our study, we are concerned with how users feel about the effect or outcome of the analysis and use of their data—whether the personalization is based on data the user provided, personal data that was explicitly inferred, or another pattern or latent inference that proves to be a proxy for personal data.

In a 2015 article in *Science* magazine, Eric Horvitz and Deirdre Mulligan discuss health-related inferences.[26] While the U.S. has additional laws to protect health data, machine learning can be

---

[23] Solon Barocas and Helen Nissenbaum. "Big Data's End Run around Anonymity and Consent" in *Privacy, Big Data, and the Public Good*, ed. Julia Lane, Victoria Stodden, Stefan Bender, Helen Nissenbaum (Cambridge University Press, 2014), 46.
[24] Cynthia Dwork and Deirdre K. Mulligan, *"*It's Not Privacy and It's Not Fair,*" Stanford Law Review Online* 35 September 3, 2013.
[25] See e.g. Federal Trade Commission, "*Data Brokers: A Call for Transparency and Accountability,*" May 2014, 47 (Personalizations may also be targeted to users based on seemingly harmless categories such as "Dog Owner" or "Health and Wellness Interest," as well as more sensitive categories like "Mobile Mixers," mostly low income Latinos and African Americans, and "Rural Everlasting," single people over 66 with limited formal education and net worth.)
[26] Eric Horvitz and Deirdre Mulligan. *"Data, privacy, and the greater good," Science*, July 17, 2015 Vol. 349(6245). 253-255.
http://research.microsoft.com/en-us/um/people/horvitz/science-2015-horvitz-mulligan-253-5.pdf

used to “infer new meaning within and across contexts and is generally unencumbered by privacy rules in the United States,” (253). Horvitz and Mulligan argue that by limiting who has access to information about people, current privacy laws effectively limit the use of personal information to discriminate. This indicates the existence of values embedded in existing data practices that, while not explicitly stated, need to be mapped over to new technologies and practices.

Empirical Work on Personalization

In a 2009 study, Turow, King, Hoofnagle, Bleakley, and Hennessy[27] looked at how users feel about targeted advertising and tracking mechanisms. They found that the majority of users (66%) did not want targeted advertising at all, and that once users are informed of the tracking mechanisms used to target ads, even more users (73% to 86%) did not want targeted advertising. These findings contradict the claims that users want targeted advertising and find it beneficial, but they are vulnerable to critiques that users say they do not want personalization yet their marketplace behaviors belie their stated preferences. The idea that people’s stated privacy preferences contradict their real world behaviors is known as the “privacy paradox.” Whether the privacy paradox arises from user choice or a lack of transparency and viable alternative choices for users who want to opt out is the subject of debate. Building on Turow et al.’s work, McDonald and Cranor conducted a study of users’ knowledge about and perceptions of online behavioural advertising.[28] They observed a group of users who wanted more relevant ads (18%), another group who saw targeted ads as creepy (46%), and a third group who said they would not notice the ads (38%).

Hannak, Soeller, Lazer, Mislove, and Wilson developed a method to detect cases of personalization in online pricing in the form of price steering (changing the order of results based on price) and price discrimination (charging different prices for different users). They define online personalization as occurring “when an inconsistency in product search results is due to [the] client-side state associated with the request” — that is, when the results are personalized based on tracking cookies, information about the user’s browser and operating system, and/or the user’s IP address.[29] The authors identify personalization in the form of price steering or discrimination on nine of the sixteen retail sites studied. Mikians, Gyarmati, Erramilli,

[27] Joseph Turow, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy, “Americans Reject Tailored Advertising and Three Activities that Enable It,” (September 29, 2009). Available at SSRN: http://ssrn.com/abstract=1478214 or http://dx.doi.org/10.2139/ssrn.1478214.
[28] Aleecia M. McDonald and Lorrie Faith Cranor, “Beliefs and behaviors: Internet users’ understanding of behavioural advertising,” *Telecommunications Policy Research Conference*, 2010.
[29] Aniko Hannak, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson. “Measuring Price Discrimination and Steering on E-commerce Web Sites,” *Proceedings of Internet Measurement Conference* (IMC 2014). Vancouver, BC, Canada, (November 2014).

and Laoutaris[30] find price differences based on location, user profile (e.g. budget conscious), and user path to the retail site (direct or through a discount aggregator).

In another related study, Hannak, Sapiezynski, Molavi Kakhki, Krishnamurthy, Lazer, Mislove, and Wilson[31] measure personalization in Google Web Search. Personalized search includes showing users locally relevant results for generic search terms like "pizza." Hannak et al. measured differences in search results for 200 users and find that on average 11.7% of their search results were different due to personalization. The authors observed the most personalization for search terms related to politics, news, and local businesses and find that measurable personalization occurs based on user location and whether users are logged in to Google.[32]

In their 2012 study, Malheiros et al. found that increasing personalization in ads increased the amount of time participants spent looking at them but also increased their discomfort with the ads. In this study, personalization was defined as "inclusion of information in the ad content that identifies or characterizes the recipient,"[33] and included the use of details from the participants' current browser sessions, as well as their names and photos. In this study the personalization was obvious to participants and the users' data appeared explicitly in the content of the ad.

A 2016 Pew Research Center study on privacy and information sharing found that Americans' opinions about information sharing were heavily context-dependent. In open-ended responses to a scenario in which a social media website provides a way to manage communications about a high school class reunion free of charge in return for tracking users' activities, people generally found it *not* acceptable (51% not acceptable, 33% acceptable, 15% it depends). In this scenario, the social media website offered a product free of charge in return for tracking user activity for later use in delivering targeted ads. However this study did find a variety of circumstances under which many Americans would share personal information in return for something of value but even in these instances (this study presented six in all) the answers people gave were shaped by the conditions of the deal.[34]

---

[30] Jakub Mikians, László Gyarmati, Vijay Erramilli, and Nikolaos Laoutaris, "Detecting price and search discrimination on the internet," *In Proceedings of the 11th ACM Workshop on Hot Topics in Networks (HotNets-XI)*, ACM, New York, NY, 2012, 79-84. DOI=http://dx.doi.org/10.1145/2390231.2390245.

[31] Aniko Hannak, Piotr Sapiezynski, Arash Molavi Kakhki, Balachander Krishnamurthy, David Lazer, Alan Mislove, and Christo Wilson, "Measuring personalization of web search," *In Proceedings of the 22nd international conference on World Wide Web (WWW '13)*, ACM, New York, NY, 527-538. DOI=http://dx.doi.org/10.1145/2488388.2488435

[32] Aniko Hannak, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson. "Measuring Price Discrimination and Steering on E-commerce Web Sites," *Proceedings of Internet Measurement Conference* (IMC 2014). Vancouver, BC, Canada, (November 2014), 527-8.

[33] Miguel Malheiros, Charlene Jennett, Snehalee Patel, Sacha Brostoff, and M. Angela Sasse, "Too Close for Comfort: A Study of the Effectiveness and Acceptability of Rich-Media Personalized Advertising," CHI 2012, 579.

[34] Pew Internet and American Life Project, "Privacy and Information Sharing," 2016

Research Questions

Online personalization's reliance on inference poses challenges for existing frameworks and regulations designed to ensure that consumers are treated fairly. As Barocas and Nissenbaum note, it is difficult to rely on the model of notice and consent when the data that is collected may seem innocuous initially but could later be used to infer more sensitive data or support seemingly unrelated personalizations. Given the central role of inferred personal information in supporting online personalization, the first question our research seeks to address is: how does inference affect user attitudes about personalization?

Inferring data about users necessarily raises questions of accuracy. When used in online personalization, inference identifies patterns in large groups of users in order to classify them into demographic, interest-based, or other categories which are used as the basis for personalization. In some cases this process may be automated to the point that the classification of the user is latent. In either case, since this approach is probabilistic and people are difficult to classify, some of the classifications, whether explicit or latent, will be inaccurate. Existing proposals in regulated areas such as credit, housing, and employment, attempt to mitigate the potential harms of inferred data by allowing people to correct inaccurate data,[35] indicating that in some cases accuracy may have an important effect on user attitudes towards personalization. This led us to a second research question: how does the accuracy of the inference affect users' attitudes about the use of inferred data in online personalization? This research question focuses on how users feel when they know whether or not the inferred data used as the basis for the personalization is accurate.

The importance of looking at privacy norms in context[36] led us to a third research question: how do the effects of inference and accuracy vary based on the domain in which the personalization occurs? We included three domains where users encounter personalization in their regular internet use: targeted advertising, filtered search results, and differential pricing. We also consider the data that is used as the basis for the personalization to be an important contextual factor that may influence user attitudes. This consideration gave rise to our last research question: how do the effects of inference and accuracy in different domains vary based on the type of data used as the basis for the personalization?

---

[35] Kate Crawford and Jason Schultz, "Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms," *Boston College Law Review*, Vol. 55, No. 93, 2014; (October 1, 2013). Available at SSRN: http://ssrn.com/abstract=2325784
[36] Helen Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life*, (Redwood, CA: Stanford University Press, 2009).

METHODS

Instrument

The data for this study was drawn from a three-part survey instrument we developed.[37] Using Qualtrics, a web-based research platform, we first presented respondents with three different vignettes.[38] In this experimental set-up, each vignette contained an online personalization scenario and after each vignette, participants were asked to respond to four separate Likert scale (5-point) rating questions as well as one open-response question. Understanding the importance of contextual factors on user attitudes, we presented vignettes in three different domains (*advertising, search results, and pricing*) and for each domain, we drew from six data types that were used as the basis for the personalization (*race*, *gender*, *household income level*, *city or town of residence*, *interests*, or *personal information*).[39] Additionally, for each data type, we presented to the respondent one of three methods by which this data was obtained or determined (provided, accurately inferred, or inaccurately inferred). The vignettes were purposefully simplified in order to isolate the key factors outlined above. The domains and data types were derived from observed real-world practices and encompass three of the main situations in which people encounter personalization online.

To test our hypotheses about the effects of inference and accuracy, and to observe contextual effects from the domain and data type, we presented respondents with one vignette in each domain randomly assigning the data type and the source/accuracy condition (provided, accurately inferred, and inaccurately inferred). The complete list of all vignettes used is provided in Table M1 (Appendix).

After the vignette section, the second part of the survey presented respondents with nine statements about potential personalization practices in order to gauge user awareness of current business practices (e.g., "Online retailers use people's location to show them different prices;" "Companies use algorithms to group people into demographic categories based on their online activity."). For each statement, respondents were asked three questions: 1) to rate how realistic they found the practice to be, 2) how frequently they had thought about it in the past, and 3) in the event that they had thought about it, how concerned they were with the practice. The third part of the survey contained additional questions relating to the sensitivity of data

---

[37] CPHS Approval #2015-12-8189, PI Coye Cheshire

[38] Kirsten E. Martin, and Helen Nissenbaum, "Measuring Privacy: An Empirical Test Using Context To Expose Confounding Variables," (December 31, 2015). Available at SSRN:http://ssrn.com/abstract=2709584 or http://dx.doi.org/10.2139/ssrn.2709584

[39] Katie Shilton and Kirsten E. Martin, "Mobile Privacy Expectations in Context," *TPRC 41: The 41st Research Conference on Communication, Information and Internet Policy*, (March 24, 2013), Available at SSRN: http://ssrn.com/abstract=2238707 or http://dx.doi.org/10.2139/ssrn.2238707.

types (including the data types used in the experimental vignettes), respondent demographics, and general technical knowledge.

Survey Participants

Prior to conducting the full-scale survey we conducted two pilot surveys, one via personal networks (N=49), and another on a limited set of Amazon Mechanical Turk ("mTurk") workers (N=17), resulting in a small number of non-substantive edits. Our final instrument consisted of 34 questions, which took approximately 15 minutes to complete. We recruited participants by creating a single HIT (Human Intelligence Task) on mTurk, on March 25, 2016.[40] We limited the HIT to Turkers living in the United States. We decided to accept both new and experienced workers since there was no theoretical basis in our study to select only experienced Turk workers, and we thought that including less experienced Turkers could possibly increase diversity along some dimensions. Participants (N=748) received $3 in compensation, paid through mTurk. See Table 1 for a breakdown of participant demographics.[41]

---

[40] Michael Buhrmester, Tracy Kwang, and Samuel D. Gosling, "Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data?" *Perspectives on Psychological Science*, vol. 6, no. 3-5, (2011).

[41] Please note that in our question about race, participants were able to check all that apply in order to more accurately reflect how they self identify. This means that the Response Ct column will not total 748, nor will the % of Total column total 100%. Additionally, we collected age by free entry, resulting in a mean age of 35.36. However, for reporting purposes, we have used age buckets from a PEW study for less granular reporting purposes. e.g. http://www.pewhispanic.org/2016/04/20/the-nations-latino-population-is-defined-by-its-youth/

Table 0: Participant Demographics

Participant Demographics

		Response Ct	% of Total	Mean	Median
Race					
	American Indian or Alaska Native	10	1.34%		
	Asian	70	9.36%		
	Black or African American	62	8.29%		
	Hispanic or Latino	50	6.68%		
	Native Hawaiian or Other Pacific Islander	3	0.40%		
	White	597	79.81%		
	Other	9	1.20%		
Age◊				35.39	33
	18-33	406	54.28%		
	34-49	235	31.42%		
	Boomer 50-68	95	12.70%		
	69 and older	11	1.47%		
Gender					
	Male	347	46.39%		
	Female	392	52.41%		
	Female-to-Male transgender	1	0.13%		
	Male-to-Female transgender	0	0.00%		
	Genderqueer	5	0.67%		
	Other	0	0.00%		
	Prefer not to answer	3	0.40%		
Household Income					
	Less than $20,000	118	15.78%		
	$20,000 - $39,999	211	28.21%		
	$40,000 - $59,999	175	23.40%		
	$60,000 - $79,999	106	14.17%		
	$80,000 - $99,999	60	8.02%		
	$100,000 +	77	10.29%		
Education Completed					
	Some or no high school education	2	0.27%		
	Finished high school	98	13.10%		
	Some undergraduate education	247	33.02%		
	Finished undergraduate education	265	35.43%		
	Some graduate-level education	48	6.42%		
	Finished graduate or other post-undergraduate professional degree	88	11.76%		
Political Affiliation					
	Republican	128	17.11%		
	Democrat	351	46.93%		
	Libertarian	28	3.74%		
	Green	7	0.94%		
	Independent	219	29.28%		
	Other	15	2.01%		

◊ A single participant's age is missing, so the Response Ct column in this section will total 747 instead of 748, and the % of Total column will not total 100%.

Dependent Variables

After each vignette, respondents were asked to rate the fairness of the scenario.[42] We measure fairness since fairness to consumers is the focus in much of the existing literature and regulatory frameworks, as well as current policy discussions.[43] Our dependent variable was measured on five-point Likert scale: Unfair (1), Somewhat Unfair (2), Neither Unfair nor Fair (3), Somewhat Fair (4), and Fair (5).

Independent Variables

Domain

To account for the possibility that any significant effects seen are due solely to the domain in which the personalization takes place, the same vignettes were presented in each of the three domains: *targeted advertising, filtered search results,* and *differential pricing*. These domains were chosen as three of the most pervasive touchpoints at which users are likely to be subject to personalization during daily Internet use and that also mostly fall outside of existing regulation. Much of what we do online involves search, commerce, or interacting with sites driven by ad revenue, so these contexts are likely to be easy for our respondents to relate to.

Data Type

To measure dependent variables across data types that may represent a range of sensitivities, each vignette contained one of six data types: *race*, *gender*, *household income leve*l, *city or town of residence*, *interests*, or *personal information*. These types represent a mix of basic sociodemographic items that any survey would capture (race, gender, household income level), an intentionally generic and broad type (personal information), and types that are very commonly used online (city or town of residence, which we use to narrow the more general category of location data, and interests). We recognize that the six data types we have chosen do not represent the full landscape, but we believe they make a good starting point.

Source

Websites often provide useful functionality and a streamlined user experience based on data that users are asked to provide directly (sometimes as a result of being logged into an account with basic details), but it is also common practice to infer data about a person, and to tailor user experience based on these assumptions. Indeed, the use of inference is prevalent in algorithmic

---

[42] Respondents were also asked to rate their *comfort level* with and the *acceptability* of the practice, but we do not analyze these dependent variables here. This paper represents one slice of the results from a larger scale study which involved data collection that may be used for multiple future papers.

[43] See, e.g., Federal Trade Commission, "*Data Brokers: A Call for Transparency and Accountability,*" May 2014; Federal Trade Commission, "Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers" (2012); White House. "Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy," (2012) .

decision-making. However, as this practice often takes place "under the hood" without including the user, our research aims to dig into the user perspective on the practice. For this reason, whether the data type was *inferred* or *provided* is varied within the vignettes.

Accuracy

Within vignettes for which the data type was inferred, an additional variable was added to the scenario: whether that inference was *accurate* or *inaccurate*. The accuracy variable was not added to vignettes in which the user provided the data. Here we assumed that users entered accurate data; inaccurate data would amount to a clerical error, intentional inaccuracy, or a change in the actual characteristic of the user over time. In the case of inference, the issue of accuracy is necessarily present, as it involves estimation. User attitudes about the accuracy of inferences can inform companies internal processes as well as policy decisions around proposals like the due process for automated decision-making proposed by Crawford and Schultz, Citron and Pasquale, and others.

RESULTS

Within each domain, we conducted an omnibus ANOVA test for each data type among the set of related conditions (provided, accurately inferred, and inaccurately inferred). Where the ANOVA was significant ($p < .05$) we performed individual t-tests between conditions accounting for multiple comparisons (Bonferroni correction).

The figures below show the mean fairness ratings for the use of each of the six data types in the three domains. Fairness is measured on a five-point Likert scale: Unfair (1), Somewhat Unfair (2), Neither Unfair nor Fair (3), Somewhat Fair (4), and Fair (5). The bars indicate the standard errors of the means.

Personal Information

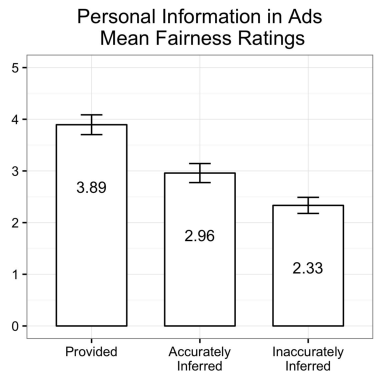

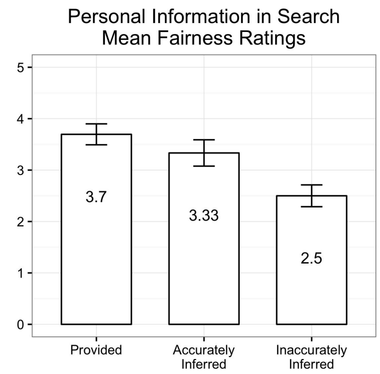

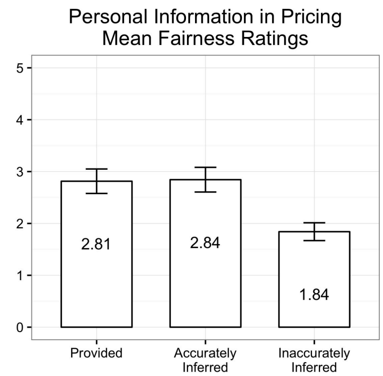

Household Income Level

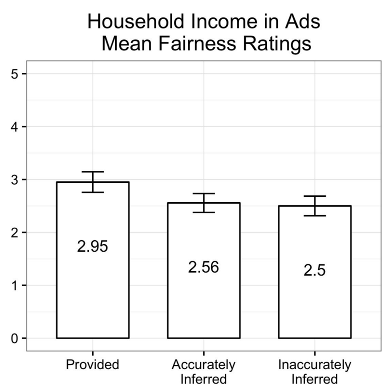

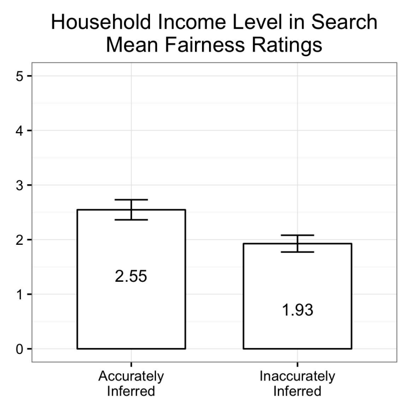

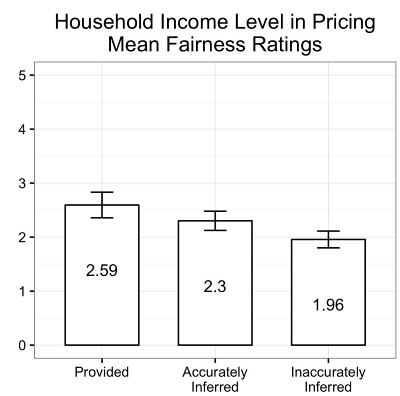

Race

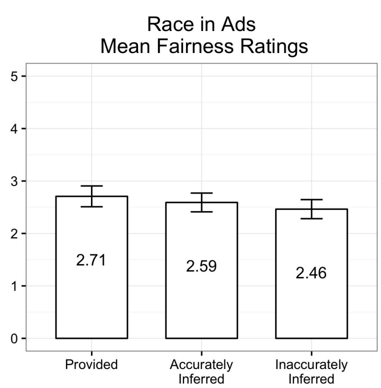

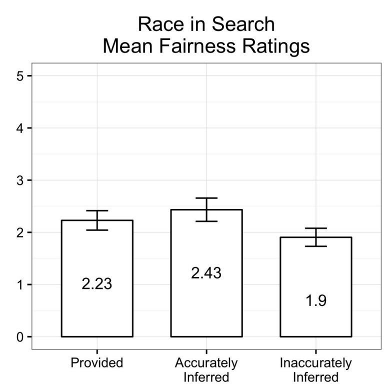

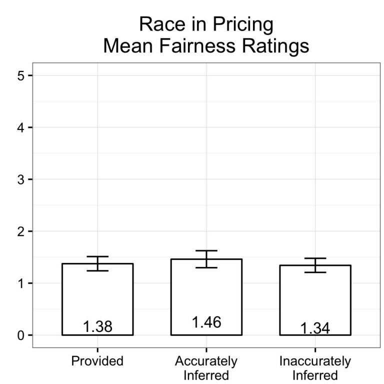

Gender

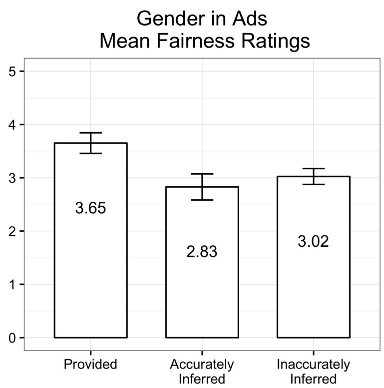

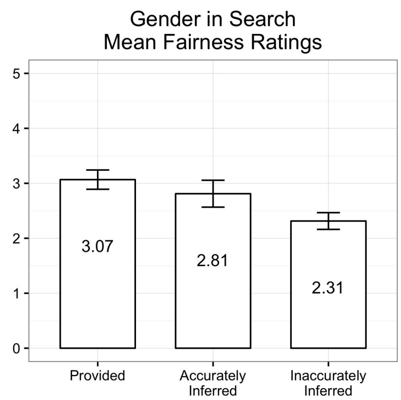

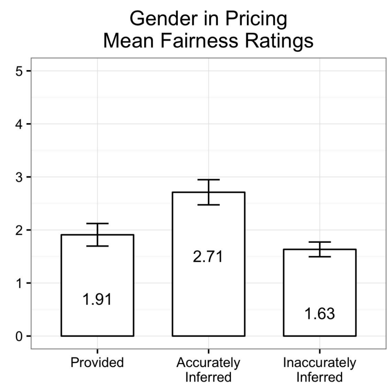

City or Town of Residence

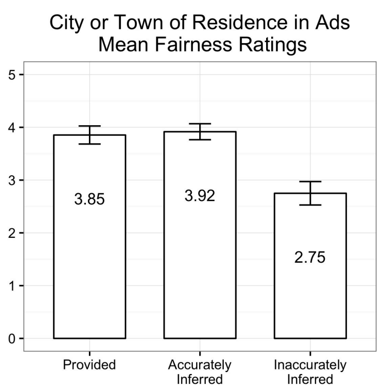

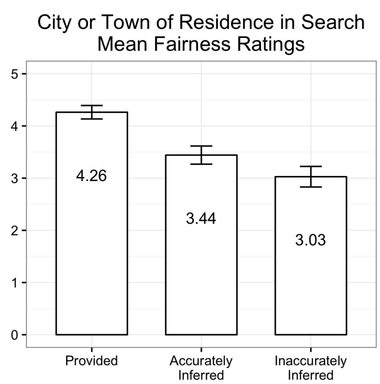

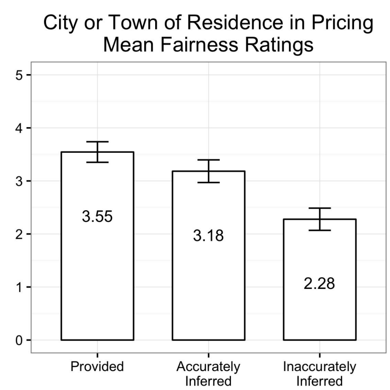

Interests

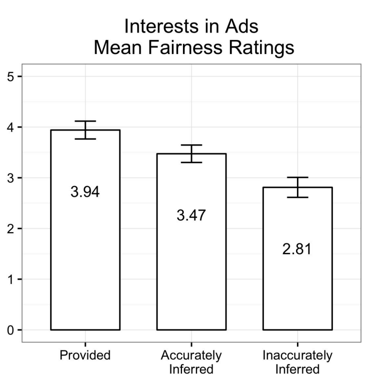

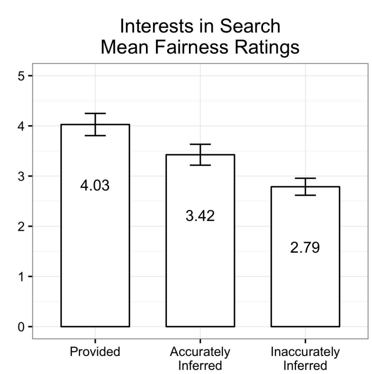

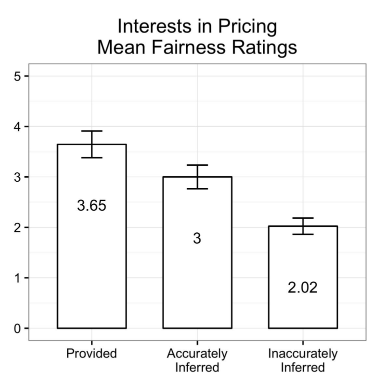

Ads

We observed that accurately inferred data has a significant effect on fairness as compared to provided data for gender ($p < .05$) and personal information ($p < .001$). For both data types we observed that the provided condition had a positive effect on fairness. The mean fairness rating for provided gender in ads was 3.65 while the mean for accurately inferred gender was 2.83 (a rating of 1 is "Unfair," a rating of 2 is "Somewhat Unfair," and 3 is "Neither Unfair nor Fair."). The mean rating for provided personal information was 3.89 and the accurately inferred mean was 2.96.

We observed a significant effect on fairness for accurately inferred data compared to inaccurately inferred data for city or town of residence ($p < .001$), interests ($p < .05$), and personal information ($p < .05$). The effects of accurately inferred data compared to inaccurately inferred data on fairness were positive for all three data types.

Search[44]

We observed a significantly higher mean fairness rating with the use of accurately inferred data than with the use of inaccurately inferred data for personal information ($p <.05$) as well as borderline significance for interests ($p = .06$). We did not find significant effects from accuracy for city or town of residence, race, or gender. Overall, the mean fairness ratings for race in search were low (mean = 1.90 for inaccurately inferred, 2.43 for accurately inferred, and 2.23 for provided). The means for gender ranged from 2.31 to 3.07.

For city or town of residence, we found a significant difference in fairness between accurately inferred and provided. The mean fairness rating was lower for accurately inferred (mean = 3.44) than for provided (mean = 4.26), representing a significant difference ($p < .001$). For interests, personal information, race, and gender we did not find significant differences in fairness between the use of accurately inferred and provided data.

Pricing

We found that the use of accurately inferred data was rated as more fair than the use of inaccurately inferred data for personal information ($p < .01$), interests ($p < .01$), city or town of residence ($p < .01$), and gender ($p < .001$). We did not observe a significant difference between accurately inferred and inaccurately inferred data for household income level or race.

The only data type for which we observed a significant difference in fairness between accurately inferred and provided data was gender ($p < .05$). For gender, we found that the use of accurately inferred data was seen as more fair than the use of provided data.

---

[44] Note: Because of an error in the survey, responses were not collected for the use of provided household income level in search so we did not run ANOVA or t-tests for household income level in search.

Data Type Sensitivity

Table 1: Data Type Sensitivity Ratings

Data Type	Mean	SD	Median
Household Income Level	3.34	0.79	4
Health Metrics (like heart rate) from a wearable device	3.29	0.95	4
Browsing History	3.23	0.89	3
Current Location	2.95	1.04	3
Purchasing Habits	2.78	0.97	3
Race and Ethnicity	2.54	1.06	3
City or Town of Residence	2.49	1.01	2
Education Level	2.44	0.95	2
Interests	2.26	0.95	2
Age	2.23	0.95	2
Gender	2.05	0.97	2

Scale: Not at all Sensitive (1), Not Too Sensitive (2), Somewhat Sensitive (3), Very Sensitive (4)

We asked respondents to rate the sensitivity of each of the data types in table 7, presenting the list as “a range of information that others might learn about you in daily life.” They were asked to “indicate how sensitive you consider that information to be (even if some people and organizations already have access to it).” The mean sensitivity ratings ranged from 2.05 for gender to 3.34 for household income level. The distributions of the responses varied more across the data types with median scores from 2 (Not Too Sensitive) to 4 (Very Sensitive). The data types that were seen as very sensitive based on the median rating are household income level and health metrics (like heart rate) from a wearable device. The data types that were seen as somewhat sensitive based on median ratings are browsing history, purchasing habits, race and ethnicity, and current location. Age, city or town of residence, education level, gender, and interests were seen as not too sensitive.

Limitations

We focused on four main factors in our vignettes: source (provided vs. inferred), accuracy, domain, and data type. In order to measure the effects of these factors we used purposefully simple vignette text. Online personalization is affected by many other factors which we did not examine. In particular, user attitudes about the sourcing of data from third parties such as data brokers merits empirical exploration, as does the purposes for which users provide their data and how its subsequent use in personalization aligns with or violates those. In this study, we did not test different factors that might affect how users feel about the use of provided data. Provided data may involve more complexity in terms of what kind of notice the user was given, and how the context of use differs from the context in which it was provided. Additionally, in our vignette text, our provided condition explicitly stated that the user provided the data to the current website and thus did not address the issue of third party data sharing.

It is important to note that the hypotheses and results in this paper are associational, not causal. This is a limitation of all cross-sectional surveys of attitudes and behaviors.

Lastly, it is possible that our use of Amazon Mechanical Turk ("mTurk") may be seen as a limitation, as some studies have found that mTurk samples are less representative than Internet-based panels or national probability samples.[45] However, here is considerable research showing that recruitment of participants from mTurk can lead to comparable or even more representative sampling of the U.S. population than other common sampling methods.[46]

DISCUSSION OF FINDINGS

This study measures how two important aspects of algorithmic personalization, the inference of personal information and the accuracy of those inferences, affect user attitudes towards online personalization in three domains: targeted ads, filtered search results, and differential pricing. We first look at the use of the generic data type "personal information" to assess users' attitudes towards personalization in each domain, and observe the effects that inference and accuracy have on perceptions of fairness. We then consider users' perceptions of fairness when a specific type of personal information (race, household income level, gender, or city or town of residence) is used as the basis for these same types of personalization. For these specific data types, we again look at the effect that inferences and the accuracy of those inferences have on user attitudes toward the personalization.

Personal Information

By positioning personal information as a generic data type, we are able to look to the fairness ratings and get a sense of how users feel about personalization in the different domains without consideration of specific data types. In doing so, we see that the use of provided personal information to target ads and filter search results is seen as neutral to somewhat fair, while the use of personal information to show users different prices is not seen as fair, regardless of whether the data is provided or inferred.

While in some cases users seem open to personalization based on provided data in ads and search, when targeted ads rely on inferred data they are seen as less fair. Using inferred

[45] Berinsky AJ, Huber GA, Lenz GS (2012) Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Analysis* 20(3): 351–368;
[46] Mullinix et al., 2015; Clifford, Scott, Ryan M Jewell, Philip D Waggoner. "Are samples drawn from Mechanical Turk valid for research on political ideology?" Research & Politics Dec 2015, 2 (4); Weinberg JD, Freese J, McElhattan D (2014) Comparing data characteristics and results of an online factorial survey between a population-based and Crowdsource-recruited sample. *Sociological Science* 1: 292–310.

personal information to filter search results may be at least tolerable to users, but it also correlates with a decrease in fairness when compared with the use of user-provided information.

Users who are subject to tailored content may not be aware of what personal information serves as the basis for the targeting. Often users are not aware that the personalization is occurring at all. However, when instances of personalization are revealed to the public there may be a negative response when the personalization differentially affected people based on membership in a demographic group, even if the effect is not intentional.[47] Looking at user attitudes towards the use of specific types of personal information in personalization can help us understand some of the norms and values that make some forms of targeting and personalization seem innovative and useful and others seem unfair and discriminatory.

Contextualizing User Attitudes About Personalization

Data Type Sensitivity

Some approaches to addressing the potential harms of inference and automated personalization attempt to identify certain types of data that should not be inferred or used in personalization and others that are acceptable. In our study we found that while there are some data types that follow this clear dichotomy, for other data types there may not be an overarching rule and instead the relevance and sensitivity of that data type must be considered in relation to the personalization it is used to support.

Avoid Personalization Based on Race and Household Income Level

Our results clearly show that personalization based on race or household income level is viewed negatively by users. While the responses to both were negative across the board, they seem to speak to different underlying issues. Household income level is seen as very sensitive by users, even when it is not being used to tailor content. This may indicate that people do not want companies to collect or infer their income in the first place. Many respondents had strong negative responses to the use of inferred household income level, raising concerns about the privacy of the data. In one case, a respondent raised suspicion about the provision of this data type: *"A website that records your personal income seems like it had to obtain in a suspicious way. I think personal income is a private data."*

Users found the use of income to target ads, filter search results,[48] and show different prices to be unfair, even when the user provided it. People did not see household income as relevant for personalization. As one respondent stated, *"I don't think that ads should be shown only based*

---

[47] See, e.g., Jennifer Valentino-Devries, Jeremy Singer-Vine, and Ashkan Soltani, "Websites Vary Prices, Deals Based on Users' Information," *Wall Street Journal*, December 24 2012.; Julia Angwin, Surya Mattu, and Jeff Larson, "The Tiger Mom Tax: Asians Are Nearly Twice as Likely to Get a Higher Price from Princeton Review," *ProPublica,* September 1, 2015.

[48] Note: We did not collect data about provided household income in search due to a surveying error. As a result we were not able to run an omnibus test on household income across all sources, and therefore excluded household income in search from our analysis.

on an individual's income level. An individual who may not have such a big income, may still desire to purchase higher priced items."

The use of race in personalization was also seen as unfair across all three domains, regardless of whether it was provided or inferred and accurate or inaccurate. However, when asked how they would feel about someone learning their race in daily life, respondents indicated that they see information about their race as not too sensitive. Users' negative perceptions of the use of race to personalize content do not necessarily arise from the sense that their race is private. Rather, for some, these perceptions arise from the view that personalization based on race does not provide relevant content: "*Prices based on race is definitely unfair and wouldn't be accurate because people of different races have varying income levels within that race."*

A large number of respondents connected the use of race to broader discrimination and raised concerns about the societal implications of showing different content to people based on race:

> *"I can see why an advertiser would try to do this, but I don't think an entire race has the same interests or characteristics. I also think it helps perpetuate racial stereotypes if ads are targeted based on race."*
>
> *"As if we needed more herding into bubbles and echo chambers, now they're going to show white people some things and black people other things? There goes our hope for common understanding -- if we aren't even getting the same results from searches."*

The use of household income level and race to personalize content was seen as unfair, regardless of the domain, the source (provided or inferred), and the accuracy. The use of these data types, or their proxies, is likely to be viewed negatively by users; ensuring accuracy of inferences or even using data that the users provided themselves does not mean that users will see them as fair.

Gender Should Be Used With Care

Respondents had mixed attitudes about the use of gender in personalization. Some were fine with the idea of companies targeting ads based on gender. As one respondent stated, *"Certain items may only (or mostly) pertain to one gender or the other, so this seems acceptable."* Another wrote, *"I don't love the idea of targeted ads in general but something based on my gender seems relatively harmless."* However, such feelings were not unanimous; some questioned the relevance of gender to ad targeting and raised concerns about exclusion: *"My gender does not reflect my likes or dislikes. I would be offended by filtering results for me. I could be missing out on something, I actually like."*

Gender was seen as not too sensitive by respondents outside of the context of personalization. Our results indicate that while users may not have an issue with companies knowing their gender, companies should be cautious when considering the use of gender (inferred or provided) as the basis for tailoring content. One respondent clearly articulated this distinction: *“I'm only uncomfortable with the fact that the prices change depending on gender?!? That's so odd! I don't care if they've inferred I'm a woman; it's the prices thing that would bug me.”*

Some forms of differential pricing based on gender map to people’s expectations, when the price differences arise because the products themselves are different. However, users found showing different prices to people based on their gender for the same products to be unfair. One respondent expressed his outrage, saying:

> *“Since I'm male, the product could be more or less expensive than the price a female would pay? That seems very unfair. Products should cost the same regardless of gender. I would be very upset if I found out that I paid more for something than a female friend on the same website on the same day. I don't trust a website that would pull something like this.”*

In search, users questioned the relevance of gender and the potential exclusionary effects. One person who expressed these concerns noted: *“I don't think my gender provides enough information to improve results, and I worry that the filtering will be based on gender stereotypes or lead to people of different genders having access to different information.”*

People did not view the use of inferred gender as fair. For ads targeted based on gender, provided gender had a significant positive effect on perceptions of fairness when compared to inferred gender. However, in search and pricing the use of provided gender does not appear to significantly improve users’ sense of fairness. When inferred gender is used to filter search results or show different prices, accuracy does have some positive effect, but not enough for accurate inferences to be seen as fair.

City or Town of Residence Is Relevant for Some Personalization

User attitudes about the use of city or town of residence to personalize content indicate that it is subject to different norms than the other types of personal information included in this study. City or town of information is viewed as not too sensitive by users and its use to personalize content was viewed as neutral or fair across the three domains.

Targeting ads based on city or town of residence was seen as reasonable and potentially useful: *“I feel the purpose of ads is advertise stuff that is meaningful to me. Getting an ad about a product or store in my town is perfectly acceptable and beneficial to me.”* Furthermore, the use of city or town of residence was not seen as discriminatory or overly personal. One person who articulated this stated: *“[T]argeted advertising based on location makes more sense and is less discriminatory than advertising based on, say, gender or race.”*

There is a sense that people are accustomed to localized search results for generic search terms (e.g. getting local coffee shop results for the search term “coffee”). As one respondent stated, *“I can see this being very helpful especially if I'm searching for concerts or things to do in my area for the weekend. Maybe I'm shopping for certain items that I hope I can locate in my area.”* In this case, the convenience provided by localized results is viewed as particularly helpful; however, some respondents raised concerns about search results based on city or town when location is not seen as relevant to their current search. One respondent stated: *“I would find this acceptable for some things. If I was searching for a book or a movie plot, I wouldn't and it wouldn't be helpful, but if I wanted the weather or a nearby restaurant, it would save time if the results were catered to where I live.”*

Even in the sensitive domain of differential pricing, people may expect or at least be able to rationalize stores charging different prices based on the varying costs of living and doing business in different cities or towns.

> *“My initial reaction is that price discrimination by geography is unfair. It doesn't seem appropriate to charge me more for living someplace. After further thought, this effectively mimics the physical world where prices in areas with higher costs and typically higher salaries tend to be more.”*

The use of city or town of residence aligns with real-world norms of location-based commerce. Users seem to want or at least to be open to tailored content based on city or town of residence when location has bearing on that particular type of content (e.g. ads or search results for local stores or events and prices that differ based on cost of living). Given the fact that this targeting can be useful to users, it is important to ensure that it is accurate. The results indicate that the use of inferred city or town of residence will not be viewed as unfair by users though when filtering search results, the use of provided data has a significant positive effect on fairness compared to the use of accurately inferred city or town of residence.

While location-based personalization at the coarser city or town level may be seen as fair to users, the sensitivity of current location indicates that more granular location-based personalization could be seen as invasive. Companies should also be conscious of the potential for location data to serve as a proxy for data types such as race and income since their use in personalization is seen very negatively.

CONCLUSION

Algorithmic personalization plays an increasing role in the way that users interact with online content. Much of this personalization is convenient for—and even expected by—users, helping them find what they want. As many in academia and government have pointed out, however, the implementation of this convenience involves collection and use of information about users, in ways that are often opaque. As we move forward into this ever more automated world, we believe that it's important to uncover these mechanisms to the users subject to the personalization, to ensure that these advances reflect user norms and expectations. Our study aims to dig into how users perceive certain elements of these practices when pieces of the process are uncovered for them. In particular, we looked at how inference and the accuracy of inference affects user attitudes toward online personalization, and how these attitudes might vary based on different types of data used, and the domain or situation in which the personalization occurs.

We found that any correlation between price and race or gender in the user's' perception will create a sense of unfairness for users. On the other hand, when the basis for the personalization is seen as relevant to the task at hand, such as basing restaurant search results on a users' city or town of residence, users are more likely to see the tailoring as fair. In these cases, users also value accuracy and participation in the form of providing data. This may indicate that data practice mechanisms are more relevant once companies and organizations have already ensured that they are not personalizing in ways that violate contextual norms.

Policy Implications

These findings speak to the need for a nuanced approach to determining if a personalization practice is likely to sit well with users. Decisions about personalization should consider the sensitivity of different types of personal information as well as their contextual relevance. This approach has the potential to be used in a variety of settings, both in industry and civil society. This kind of norms-based approach to inference has been articulated by privacy scholars[49] as "contextual integrity," "cross-context use constraints," and limiting "context-jumping." The idea of referring to contextual norms as a guideline for limits on inference has been proposed by Solon Barocas, who describes it as "inferential privacy."[50] Specifically, he argues that "decisions based on contextually distant proxies are likely to provoke objections as violations of privacy because they involve inferences that draw on criteria that seem unrelated to the things with which they nevertheless correlate."[51]

One approach others have proposed to uncover these cultural norms is to say that a company should not infer individual characteristics from data if it would not be asked in person. For example, if it would be inappropriate to ask about a person's income level when they're standing in front of you, it is similarly inappropriate to infer this information about a user and to

[49] Helen Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life*, (Redwood, CA: Stanford University Press, 2009).; Scott R. Peppet, "Regulating the Internet of Things: First Steps Toward Managing Discrimination, Privacy, Security, and Consent," *Texas Law Review* Vol. 93:85, 2014.; Barocas; Horvitz and Mulligan
[50] Barocas, Solon "Leaps and Bounds" Talk at Princeton University November 2015. https://www.youtube.com/watch?v=d8jdgxJHW4Q
[51] Ibid.

personalize their online experience based on the insight.[52] While this approach is useful when the inferred information is sensitive, Barocas argues that it does not actually tell us what the norm is. Our findings suggest that Barocas's argument for contextual integrity in the process of inferring information could be extended to personalizing content. An approach to personalization that respects contextual integrity would draw on cultural norms to avoid personalizing based on seemingly unrelated characteristics. Creating a rubric focused around cultural norms of contextual integrity requires an understanding of user attitudes. The findings from this research can guide companies in implementing practices that are grounded in existing frameworks like the FIPPs and also take the importance of a user-centered understanding of context into account.

Future Work

The results of this study indicate that users have nuanced and context-dependent attitudes about personalization that we believe merit further exploration through additional surveys and vignette studies, controlled experiments, and interviews. We would also like to see the focus of our work examined through research methods from the human-computer interaction (HCI) field such as those employed in work by Eslami et al.,[53] wherein contextual inquiry, interview, and survey methods were used to develop an understanding of user awareness of and attitudes toward the Facebook News Feed algorithm.

Future research could also work to tease apart the nuances of the results found here. For example, building on our findings regarding the use of city or town of residence to personalize, future studies could look at more granular location data to understand when location-based personalization violates user expectations and norms.

ACKNOWLEDGMENTS

We would like to thank the Center for Technology, Society & Policy, and the Center for Long-Term Cybersecurity for their funding of this project. The research team would like to thank our advisor Coye Cheshire for his excellent guidance on our research design and analysis. We would also like to thank Chris Hoofnagle, Jen King, and Richmond Wong for their thoughtful comments.

---

[52] See e.g. Chris Hoofnagle, Federal Trade Commission Privacy Law and Policy (2016)(discussing retailer use of reverse-zip-code lookups to determine customers home addresses because asking customers indirectly avoided "losing customers who feel that you're invading their privacy.").

[53] Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, Christian Sandvig. "I always assumed that I wasn't really that close to [her]": Reasoning about invisible algorithms in the news feed. 2015 http://www-personal.umich.edu/~csandvig/research/Eslami_Algorithms_CHI15.pdf

APPENDIX A - VIGNETTES

Table M1: Vignette Universe							
TARGETED ADVERTISING							
CONTEXT			DATA TYPE		SOURCE		ACCURACY
You are reading an article on a website.	An ad is shown to you	based on	your race	which	you provided	to this website.	
You are reading an article on a website.	An ad is shown to you	based on	your gender	which	you provided	to this website.	
You are reading an article on a website.	An ad is shown to you	based on	your household income level	which	you provided	to this website.	
You are reading an article on a website.	An ad is shown to you	based on	your city or town of residence	which	you provided	to this website.	
You are reading an article on a website.	An ad is shown to you	based on	your interests	which	you provided	to this website.	
You are reading an article on a website.	An ad is shown to you	based on	your personal information	which	you provided	to this website.	
You are reading an article on a website.	An ad is shown to you	based on	your race	which	was inferred	from the webpages you visit	and is accurate.
You are reading an article on a website.	An ad is shown to you	based on	your gender	which	was inferred	from the webpages you visit	and is accurate.
You are reading an article on a website.	An ad is shown to you	based on	your household income level	which	was inferred	from the webpages you visit	and is accurate.
You are reading an article on a website.	An ad is shown to you	based on	your city or town of residence	which	was inferred	from the webpages you visit	and is accurate.
You are reading an article on a website.	An ad is shown to you	based on	your interests	which	were inferred	from the webpages you visit	and are accurate.
You are reading an article on a website.	An ad is shown to you	based on	your personal information	which	was inferred	from the webpages you visit	and is accurate.
You are reading an article on a website.	An ad is shown to you	based on	your race	which	was inferred	from the webpages you visit	and is inaccurate.
You are reading an article on a website.	An ad is shown to you	based on	your gender	which	was inferred	from the webpages you visit	and is inaccurate.
You are reading an article on a website.	An ad is shown to you	based on	your household income level	which	was inferred	from the webpages you visit	and is inaccurate.
You are reading an article on a website.	An ad is shown to you	based on	your city or town of residence	which	was inferred	from the webpages you visit	and is inaccurate.
You are reading an article on a website.	An ad is shown to you	based on	your interests	which	were inferred	from the webpages you visit	and are inaccurate.
You are reading an article on a website.	An ad is shown to you	based on	your personal information	which	was inferred	from the webpages you visit	and is inaccurate.
FILTERED SEARCH RESULTS							
CONTEXT			DATA TYPE		SOURCE		ACCURACY
You are using a search engine.	Your search results are filtered	based on	your race	which	you provided	to this search engine.	
You are using a search engine.	Your search results are filtered	based on	your gender	which	you provided	to this search engine.	

You are using a search engine.	Your search results are filtered	based on	your city or town of residence	which	you provided	to this search engine.	
You are using a search engine.	Your search results are filtered	based on	your interests	which	you provided	to this search engine.	
You are using a search engine.	Your search results are filtered	based on	your personal information	which	you provided	to this search engine.	
You are using a search engine.	Your search results are filtered	based on	your race	which	was inferred	from the webpages you visit	and is accurate.
You are using a search engine.	Your search results are filtered	based on	your gender	which	was inferred	from the webpages you visit	and is accurate.
You are using a search engine.	Your search results are filtered	based on	your household income level	which	was inferred	from the webpages you visit	and is accurate.
You are using a search engine.	Your search results are filtered	based on	your city or town of residence	which	was inferred	from the webpages you visit	and is accurate.
You are using a search engine.	Your search results are filtered	based on	your interests	which	were inferred	from the webpages you visit	and are accurate.
You are using a search engine.	Your search results are filtered	based on	your personal information	which	was inferred	from the webpages you visit	and is accurate.
You are using a search engine.	Your search results are filtered	based on	your race	which	was inferred	from the webpages you visit	and is inaccurate.
You are using a search engine.	Your search results are filtered	based on	your gender	which	was inferred	from the webpages you visit	and is inaccurate.
You are using a search engine.	Your search results are filtered	based on	your household income level	which	was inferred	from the webpages you visit	and is inaccurate.
You are using a search engine.	Your search results are filtered	based on	your city or town of residence	which	was inferred	from the webpages you visit	and is inaccurate.
You are using a search engine.	Your search results are filtered	based on	your interests	which	were inferred	from the webpages you visit	and are inaccurate.
You are using a search engine.	Your search results are filtered	based on	your personal information	which	was inferred	from the webpages you visit	and is inaccurate.
DIFFERENTIAL PRICING							
CONTEXT			DATA TYPE		SOURCE		ACCURACY
You are shopping on a retail website.	The prices that you see are	based on	your race	which	you provided	to this retailer.	
You are shopping on a retail website.	The prices that you see are	based on	your gender	which	you provided	to this retailer.	
You are shopping on a retail website.	The prices that you see are	based on	your household income level	which	you provided	to this retailer.	
You are shopping on a retail website.	The prices that you see are	based on	your city or town of residence	which	you provided	to this retailer.	

You are shopping on a retail website.	The prices that you see are	based on	your interests	which	you provided	to this retailer.	
You are shopping on a retail website.	The prices that you see are	based on	your personal information	which	you provided	to this retailer.	
You are shopping on a retail website.	The prices that you see are	based on	your race	which	was inferred	from the webpages you visit	and is accurate.
You are shopping on a retail website.	The prices that you see are	based on	your gender	which	was inferred	from the webpages you visit	and is accurate.
You are shopping on a retail website.	The prices that you see are	based on	your household income level	which	was inferred	from the webpages you visit	and is accurate.
You are shopping on a retail website.	The prices that you see are	based on	your city or town of residence	which	was inferred	from the webpages you visit	and is accurate.
You are shopping on a retail website.	The prices that you see are	based on	your interests	which	were inferred	from the webpages you visit	and are accurate.
You are shopping on a retail website.	The prices that you see are	based on	your personal information	which	was inferred	from the webpages you visit	and is accurate.
You are shopping on a retail website.	The prices that you see are	based on	your race	which	was inferred	from the webpages you visit	and is inaccurate.
You are shopping on a retail website.	The prices that you see are	based on	your gender	which	was inferred	from the webpages you visit	and is inaccurate.
You are shopping on a retail website.	The prices that you see are	based on	your household income level	which	was inferred	from the webpages you visit	and is inaccurate.
You are shopping on a retail website.	The prices that you see are	based on	your city or town of residence	which	was inferred	from the webpages you visit	and is inaccurate.
You are shopping on a retail website.	The prices that you see are	based on	your interests	which	were inferred	from the webpages you visit	and are inaccurate.
You are shopping on a retail website.	The prices that you see are	based on	your personal information	which	was inferred	from the webpages you visit	and is inaccurate.

APPENDIX B - RESULT TABLES

Table 2: Significant Results in Ads

	Condition: Provided/Accurately Inferred	
	t-stat	*Cohen's d*
Context: Ads		
Dependent Variable: Fairness		
Gender	-2.639*	0.605
Personal Information	-3.519***	0.756
Dependent Variable: Trustworthiness		
Gender	-2.485*	0.568

	Condition: Accurately Inferred/Inaccurately Inferred	
	t-stat	*Cohen's d*
Context: Ads		
Dependent Variable: Fairness		
Personal Information	2.594*	0.526
Dependent Variable: Trustworthiness		
Personal Information	2.848*	0.578

*p < 0.05, **p < 0.01, ***p < 0.001, †p ≤ 0.1.

Table 3: Descriptive Results in Ads

	Provided	Accurately Inferred	Inaccurately Inferred
	M (SD)	M (SD)	M (SD)
Context: Ads			
Dependent Variable: Fairness			
City or Town of Residence	3.854 (1.18)	3.917 (1.05)	2.75 (1.17)
Gender	3.651 (1.27)	2.829 (1.44)	3.024 (0.96)
Household Income Level	2.951 (1.24)	2.556 (1.20)	2.500 (1.05)
Interests	3.941 (1.26)	3.477 (1.06)	2.811 (1.20)
Personal Information	3.895 (1.18)	2.960 (1.29)	2.333 (1.08)
Race	2.707 (1.27)	2.591 (1.19)	2.463 (1.16)
Dependent Variable: Trustworthiness			
City or Town of Residence	3.604 (1.16)	3.500 (1.11)	2.679 (1.28)
Gender	3.302 (1.08)	2.657 (1.19)	2.439 (0.92)
Household Income Level	2.561 (1.14)	2.356 (1.11)	2.156 (0.95)
Interests	3.549 (1.17)	3.210 (0.93)	2.649 (1.25)
Personal Information	3.237 (1.17)	2.714 (1.22)	2.083 (0.94)
Race	2.463 (1.16)	2.364 (1.10)	2.073 (1.08)

Note:
For the fairness rating questions, a larger number means that the respondent thought the practice more fair. Here the scale ranged from 1 to 5, where a rating of 1 was “Unfair,” a rating of 2 was “Somewhat Unfair,” 3 was “Neither Unfair nor Fair,” 4 was "Somewhat Fair," and 5 was "Fair."

Similarly, for the trustworthiness rating questions, the scale ranged from 1 to 5, where a rating of 1 was “Untrustworthy,” a rating of 2 was “Somewhat Untrustworthy,” 3 was “Neither Untrustworthy nor Trustworthy,” 4 was "Somewhat Trustworthy," and 5 was " Trustworthy."

Table 4: Significant Results in Search

	Condition: Provided/Accurately Inferred	
	t-stat	*Cohen's d*
Context: Search		
Dependent Variable: Fairness		
City or Town of Residence	-3.794***	0.836
Dependent Variable: Trustworthiness		
City or Town of Residence	-4.035***	0.891
Interests	-2.635*	0.601

	Condition: Accurately Inferred/Inaccurately Inferred	
	t-stat	*Cohen's d*
Context: Search		
Dependent Variable: Fairness		
Personal Information	2.507*	0.649
Dependent Variable: Trustworthiness		
Gender	2.762*	0.606
Interests	3.555***	0.766
Personal Information	3.231***	0.841
Race	2.799*	0.674

*p < 0.05, **p < 0.01, ***p < 0.001, †p ≤ 0.1.

Table 5: Descriptive Results in Search

	Provided	Accurately Inferred	Inaccurately Inferred
	M (SD)	M (SD)	M (SD)
Context: Search			
Dependent Variable: Fairness			
City or Town of Residence	4.263 (0.79)	3.442 (1.14)	3.028 (1.18)
Gender	3.067 (1.18)	2.811 (1.49)	2.314 (1.09)
Household Income Level	n/a*	2.547 (1.34)	1.927 (0.98)
Interests	4.027 (1.34)	3.425 (1.32)	2.787 (1.16)
Personal Information	3.696 (1.38)	3.333 (1.33)	2.500 (1.24)
Race	2.229 (1.29)	2.433 (1.22)	1.905 (1.12)
Dependent Variable: Trustworthiness			
City or Town of Residence	4.079 (0.82)	3.233 (1.07)	2.806 (1.19)
Gender	3.111 (1.17)	2.838 (1.34)	2.098 (1.08)
Household Income Level	n/a*	2.528 (1.20)	1.878 (0.93)
Interests	3.838 (1.19)	3.125 (1.18)	2.23 (1.15)
Personal Information	3.326 (1.30)	3.148 (1.32)	2.118 (1.12)
Race	2.230 (1.32)	2.400 (1.10)	1.690 (1.00)

**missing value*

Note:

For the fairness rating questions, a larger number means that the respondent thought the practice more fair. Here the scale ranged from 1 to 5, where a rating of 1 was “Unfair,” a rating of 2 was “Somewhat Unfair,” 3 was “Neither Unfair nor Fair,” 4 was "Somewhat Fair," and 5 was "Fair."

Similarly, for the trustworthiness rating questions, the scale ranged from 1 to 5, where a rating of 1 was “Untrustworthy,” a rating of 2 was “Somewhat Untrustworthy,” 3 was “Neither Untrustworthy nor Trustworthy,” 4 was "Somewhat Trustworthy," and 5 was " Trustworthy."

Table 6: Significant Results in Pricing

		Condition: Provided/Accurately Inferred	
		t-stat	*Cohen's d*
Context: Pricing			
	Dependent Variable: Fairness		
	Gender	2.518*	0.534
	Dependent Variable: Trustworthiness		
	Gender	2.278†	0.483
		Condition: Accurately Inferred/Inaccurately Inferred	
		t-stat	*Cohen's d*
Context: Pricing			
	Dependent Variable: Fairness		
	Gender	3.922***	0.836
	Interests	3.416***	0.784
	Personal Information	3.416***	0.804
	Dependent Variable: Trustworthiness		
	City or Town of Residence	3.139***	0.686
	Gender	3.562***	0.760
	Interests	3.76***	0.861
	Personal Information	4.30E+00	1.009

*p < 0.05, **p < 0.01, ***p < 0.001, †p ≤ 0.1.

Table 7: Descriptive Results in Pricing

	Provided	Accurately Inferred	Inaccurately Inferred
	M (SD)	M (SD)	M (SD)
Context: Pricing			
Dependent Variable: Fairness			
City or Town of Residence	3.545 (1.44)	3.184 (1.50)	2.278 (1.26)
Gender	1.909 (1.41)	2.711 (1.59)	1.634 (0.89)
Household Income Level	2.595 (1.44)	2.302 (1.29)	1.957 (1.05)
Interests	3.645 (1.47)	3.000 (1.41)	2.024 (1.05)
Personal Information	2.814 (1.55)	2.844 (1.35)	1.841 (1.14)
Race	1.375 (0.87)	1.462 (1.02)	1.343 (0.80)
Dependent Variable: Trustworthiness			
City or Town of Residence	3.382 (1.22)	3.082 (1.37)	2.167 (1.30)
Gender	2.023 (1.28)	2.667 (1.38)	1.805 (0.81)
Household Income Level	2.432 (1.26)	2.189 (1.18)	1.717 (0.78)
Interests	3.226 (1.36)	2.972 (1.28)	1.976 (1.02)
Personal Information	2.651 (1.46)	2.781 (1.13)	1.705 (1.00)
Race	1.525 (1.01)	1.744 (1.04)	1.314 (0.76)

Note:
For the fairness rating questions, a larger number means that the respondent thought the practice more fair. Here the scale ranged from 1 to 5, where a rating of 1 was “Unfair,” a rating of 2 was “Somewhat Unfair,” 3 was “Neither Unfair nor Fair,” 4 was "Somewhat Fair," and 5 was "Fair."

Similarly, for the trustworthiness rating questions, the scale ranged from 1 to 5, where a rating of 1 was “Untrustworthy,” a rating of 2 was “Somewhat Untrustworthy,” 3 was “Neither Untrustworthy nor Trustworthy,” 4 was "Somewhat Trustworthy," and 5 was " Trustworthy."