



Field Evidence of the Effects of Privacy, Data Transparency, and Pro-social Appeals on COVID-19 App Attractiveness

Samuel Dooley
sdooley1@cs.umd.edu
University of Maryland
College Park, MD, USA

John P. Dickerson
john@cs.umd.edu
University of Maryland
College Park, MD, USA

Dana Turjeman
dana.turjeman@idc.ac.il
Reichman University
Herzliya, Israel

Elissa M. Redmiles
eredmiles@mpi-sws.org
Max Plank Institute for Software Systems
Saarbrücken, Germany

ABSTRACT

COVID-19 exposure-notification apps have struggled to gain adoption. Existing literature posits as potential causes of this low adoption: privacy concerns, insufficient data transparency, and the type of appeal – collective- vs. individual-good – used to frame the app. As policy guidance suggests using tailored advertising to evaluate the effects of these factors, we present the first field study of COVID-19 contact tracing apps with a randomized, control trial of 14 different advertisements for CovidDefense, Louisiana’s COVID-19 exposure-notification app. We find that all three hypothesized factors – privacy, data transparency, and appeals framing – relate to app adoption, even when controlling for age, gender, and community density. Our results offer (1) the first field evidence supporting the use of collective-good appeals, (2) nuanced findings regarding the efficacy of data and privacy transparency, the effects of which are moderated by appeal framing and potential users’ demographics, and (3) field-evidence-based guidance for future efforts to encourage pro-social health technology adoption.

CCS CONCEPTS

• **Security and privacy** → **Social aspects of security and privacy**; *Usability in security and privacy*; • **Human-centered computing** → Empirical studies in collaborative and social computing.

KEYWORDS

COVID-19, Field Study, Pro-Social, Privacy, Transparency

ACM Reference Format:

Samuel Dooley, Dana Turjeman, John P. Dickerson, and Elissa M. Redmiles. 2022. Field Evidence of the Effects of Privacy, Data Transparency, and Pro-social Appeals on COVID-19 App Attractiveness. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*, April 30 - May 6 2022, New Orleans, LA. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3491102.3501869>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '22, April 30 - May 6 2022, New Orleans, LA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9157-3/22/04...\$15.00
<https://doi.org/10.1145/3491102.3501869>

1 INTRODUCTION

To combat SARS-CoV-2 – also known as “coronavirus” – and its associated illness COVID-19, countries and other entities have worked to develop vaccines and a variety of other mitigation tools. One such tool is contact-tracing technology that serves as the foundation for exposure-notification apps (COVID-19 apps) that can alert users when they have been exposed to coronavirus. These apps have been developed and deployed in 77 countries and U.S. states.¹

Similar to other pro-social COVID-19 behaviors such as vaccination and mask adoption, greater adoption of COVID-19 apps improves their efficacy. Yet, adoption has been low, with the highest adoption rates per jurisdiction hovering around 30% and typical adoption rates closer to 10%.²

Prior work has sought to understand people’s considerations for adopting such apps through self-report and lab-based studies. These works suggest that people’s adoption is likely driven by concerns regarding the app’s privacy and data collection practices, as well as perceptions of whether the benefits of the apps – to themselves or to society – outweigh their privacy and data concerns [43, 58, 63].

However, none of this prior work observed how these considerations affect actual adoption of these apps in the wild. Self-report studies on privacy-related behavior – such as the adoption of COVID-19 apps – have known flaws due to the “Privacy Paradox” which shows a disconnect between one’s beliefs about privacy choices and one’s actual behavior. While the evidence for this paradox can be caused by a poor translation from beliefs to behaviors, this can also be caused by measurement device biases. People will claim that they *would* choose a more privacy-oriented product or would not be willing to share information, but when observed making such decisions in real life, they often forgo privacy protections [2, 41, 52]. Thus, to offer real-world insights into users’ adoption of a privacy-sensitive health application in the context of COVID-19, we conduct the first, to our knowledge, field study of COVID-19 app adoption.

We collaborated with the state of Louisiana to conduct a randomized, controlled field experiment on the impact of tailored messaging addressing the attributes found most relevant to app adoption

¹See the Linux Public Health Foundation dashboard (<https://landscape.lfph.io/>) for a running list of deployed COVID-19 apps.

²There has been little official reporting of COVID-19 app adoption rates outside of the popular press; we refer to <https://time.com/5905772/covid-19-contact-tracing-apps/> for these adoption statistics.

in prior work – the app’s benefits, privacy, and data collection – on adoption of the state’s COVID-19 exposure-notification app, CovidDefense. Specifically, we test the impact of advertisements that contain two types of messaging addressing factors identified in prior work and recommended in policy guidance [49]: (a) app benefits framed as either a collective- or individual-good and (b) transparency regarding privacy and/or data collection.

We conducted our field experiment on the Google Ads Platform using 14 different ads. Ads were randomly displayed to Louisiana residents and generated 7,010,271 impressions.³

The outcome measured was whether the user clicked the respective ad; those who clicked were redirected to the Louisiana Department of Public Health app download page (<http://coviddefensela.com/>).⁴

Using these data we address four research questions:

- RQ1:** Is messaging that presents the benefit of app installation as a *collective-good* appeal (i.e., with societal benefit) more effective than messaging that appeals to *individual-good*?
- RQ2:** Is messaging that makes *privacy transparent* more effective than messaging that does not? And, which privacy-transparency statements are most/least effective, those that: (a) broadly reassure people about privacy concerns, or those that specifically focus on enhanced control over data collection – through a statement emphasizing either (b) general, non-technical privacy control or (c) technically concrete privacy control?
- RQ3:** Is messaging that makes *data collection transparent* (i.e., stating clearly what data is being collected by the app) more effective than messaging that does not inform potential users what data the app will collect?
- RQ4:** How do demographics (age, gender, geography) moderate the adoption of CovidDefense and the experimental effects observed in RQs1-3?

Collective-good appeals (i.e., pro-social messages that speak to community benefit) are suggested as a best practice by existing policy guidance [49]. However, the efficacy of such appeals is empirically debated in the context of COVID-19 [55] on the basis of evidence from self-report data, laboratory experiments, and hybrid self-report tracking [38, 48, 60] and the impact of these appeals in other privacy-sensitive technology settings has not been well studied.

Existing policy guidance also encourages transparency in advertising promoting pro-social health behaviors, and prior work on people’s intent to adopt COVID-19 apps emphasizes the importance of privacy and data collection concerns on people’s adoption intent [43, 58, 63, 76]. However, there is little field evidence regarding how individuals respond to data transparency and privacy statements in a privacy-sensitive health technology context. While prior research in the privacy domain (e.g., 11, 24, 68) has found that increased transparency and sense of control regarding existing privacy and data collection may reduce concerns and increase

willingness to share data, it is an open question whether such transparency can be effectively provided through tailored messaging [59]. Findings from some prior work [36] suggest that it can; but this and other work suggest that increasing the salience of privacy through transparency at the time of the choice to adopt the app could also artificially increase people’s concerns about privacy [36, 59].

The results of our field study show that tailored messaging can effectively influence the pro-social behavior of installing a COVID-19 app. We find that significantly more people click on messages that use collective-good appeals than those that use individual-good appeals (RQ1). Furthermore, in a series of moderation analyses, we find that transparency about privacy (RQ2) and data collection (RQ3) moderate this effect. Specifically, collective-goods appeals are even more effective when paired with a privacy-transparency statement, but are less effective when additionally paired with a data-transparency statement. Individual-goods appeals exhibit the opposite effects. Such differences suggest that digital privacy and data transparency can be effectively provided through tailored messaging, but we must think carefully about how an application’s purpose and framing may impact people’s privacy considerations and reasoning. Finally, our results shed light on how priming with an individual-good appeal increases gender and age differences in receptiveness to the app and to the different privacy controls presented (RQ4).

Our findings offer insight into how users make privacy-benefits trade-offs when making decisions to adopt an app in the wild. We confirm in the field prior self-report results on the importance of individualist vs. collectivist mindset [60], and expand the existing body of literature to provide insight into the real world impact of the tension between our desire to improve community health by sharing personal data and our individual desire for privacy.

2 RELATED WORK

Here, we review the prior work most closely related to our study: on the factors that influence COVID-19 app adoption and on privacy- and data-transparency statements in the context of digital health.

2.1 Factors Influencing Intent to Adopt COVID-19 Apps

When considering whether to use an app for COVID-19 contact tracing, previous research has shown three main considerations: the functionality of the app, concerns regarding privacy, and concerns regarding data collection.

The two main functions of COVID-19 contact tracing apps are to indicate to a user if they have been exposed (an individual-good) and to help the broader community reduce the spread of the virus (a collective-good) [58]. Li et al. [43] found that of these dual purposes of the app were more influential in determining intention to install than security or privacy concerns. Williams et al. [74] finds that even the possibility of a collective-good outcome can convince otherwise hesitant users to participate in COVID-19 apps, sometimes begrudgingly. However, some individuals indicate a reluctance to install a COVID-19 app regardless of how well the app works [32, 43, 63].

User privacy is well-documented as a main source of hesitancy for individuals to download and use COVID-19 apps, stemming

³As is typical in digital marketing campaigns, ads may be displayed during Google search more than once to the same user/IP address and thus the number of impressions is larger than the population of Louisiana.

⁴A user may have seen more than one ad because Google does not allow us to control this. However, if the user clicked on an ad, the click was associated with the specific ad on which they clicked.

from both general privacy concern as well as specific concerns about the privacy of contact-tracing apps [23, 32, 40, 43, 63, 74, 76, 78]. However, there is still debate about whether people's stated privacy concerns actually influence COVID-19 app adoption once controlling for other factors such as incentives [21], institutional trust [28, 31, 69], political ideology [44], and general perceptions of COVID-19 [14, 21, 43, 69, 71]. Additionally, there is indication that privacy concerns can be linked to app functionality. For instance, there is evidence that people's privacy considerations about COVID-19 apps can be moderated by the way in which the app works, specifically the centralization of the contact-tracing mechanism. There is not a consensus in prior work regarding whether users prefer a centralized or decentralized system: Zhang et al. [76] find from their conjoint analysis that a decentralized system had higher app adoption whereas other studies find the opposite [28, 43].

Another implementation choice regards the data used for contact tracing. Some COVID-19 contact-tracing apps operate using only proximity data, relying on Bluetooth to detect proximity between devices, while other apps rely on GPS location data. Prior work [58, 63] finds that users worry about data collection in general, regardless of data type, and also find that users are more comfortable with apps that use proximity vs. location data. These concerns and considerations interplay with users' concerns about their privacy, as some of these concerns focus on the privacy of the data collected by the app, even if it is stored only on their device as is the case for decentralized apps.⁵ Regardless, the majority of deployed contact-tracing apps are decentralized and use proximity data [20].

Prior field work on COVID-19 App Adoption. Due to the emerging nature of the pandemic, there has been little field work studying how people's adoption considerations influence their behavior in the real world. Our work seeks to build on findings from prior self-report work while filling the gap of empirical field evidence.

Most closely related to our work, Munzert et al. [48] tested the effect of presenting collective-good appeals in combination with privacy and functionality-related information in a video intervention on people's adoption of a COVID-19 app. Subjects in their study participated in an opt-in survey panel in Germany and their digital behavior could be tracked by the survey panel, allowing the researchers to observe whether participants installed Germany's COVID-19 app at some point after seeing the video intervention in the survey. Their experiment found no effects from the video intervention, though this might be an artifact of some experimental limitations, as identified in Toussaert [66], which include 1. the nature of the intervention, which involved exposure to a training video during a survey-based study rather than as part of real-world installation behavior and which combined multiple experimental messages, preventing isolation of the impact of the collective-good appeal from the other experimental factors, 2. the sample size, and 3. the opt-in nature of the participant pool. In contrast, our work isolates and focuses specifically on the impact of appeals in tailored messaging, presenting the first, to our knowledge⁶, direct field

evaluation in the general population of the efficacy of collective-good appeals in encouraging pro-social COVID-19 behavior, as well as the first direct field evidence of the impacts of privacy and data transparency on adoption of a pro-social health technology. Specifically, we tested this in tailored advertising messaging used to encourage adoption of an exposure notification app at the time it was released to the population. Importantly, our work does not rely on surveys, online studies or an opt-in sample. Instead, we directly measure the outcome of interest: whether a prospective user clicks to download the app.

2.2 Privacy and Digital Health

Outside of a COVID-19 setting, privacy concerns in digital and mobile health applications (often termed mHealth apps) have an extensive research history. Contact-tracing mHealth apps have been studied before in other settings such as tuberculosis [9], influenza [25], and H1N1 [62], though much of the framing of these studies has been from the perspective of the individual benefits they provide [64], and an emphasis on privacy in these studies is generally lacking given the nascentness of mHealth at the time of their study.

More broadly, privacy influences an individual's interest in adopting mHealth applications [5, 22, 53, 54]. Early on in the study of privacy effects on mHealth app adoption, Klasnja et al. [37] explored the privacy perspectives of different data types and found that GPS location data was particularly sensitive. Prasad et al. [53] found that an individual may have differing attitudes towards sharing the same data with different individuals, reporting that individuals who wore fitness trackers were less likely to share data with friends and family than with strangers. Among other things, Demographics such as age, education, occupation, and digital penetration in the user's country may also play a role in privacy perceptions around mHealth applications [33, 61].

Prior work typically finds that the main predictors of mHealth app adoption are trust, utility, and ease of use [3, 16, 17, 51, 65, 70, 77] with moderating effects from age, gender, location, and education. These findings align well with the literature on adoption intention for mHealth apps specifically designed for COVID-19 as described above.

Focusing specifically on privacy and sensitivity around data collection, Jacobs et al. [30] explores the data sharing preferences of different groups of people by role in a data ecosystem around breast cancer. They find that patients, doctors, and navigators have different comfort levels with data sharing, e.g., patients are hesitant to share data about their emotional state. Warner et al. [72] explore the specific privacy concerns within group of HIV-positive men using a geo-social dating app and find that some users disclose their status to reduce their exposure to stigma while others avoid disclosure to avoid being stigmatized.

More broadly, there is limited prior work on the effect of privacy and sense of control on the sharing of personal health related data, perhaps because individuals do not have much control over their own health data, and sometimes are not even able to access it themselves. While HIPAA and other health-related privacy policies have been developed to let users exercise informed consent over

⁵Note that it is possible for even decentralized apps to have privacy leaks [10, 32, 50, 56], and thus user's privacy concerns are not unfounded.

⁶Banker and Park conducted a field study on the impact of collective-good appeals on clicks to CDC guidelines at the very beginning of the pandemic [7]. However, health information consumption and pro-social health behavior are importantly different constructs.

sharing health information, such mechanisms are dated and may not be applicable for mHealth [47].

Outside of the health domain, prior work on privacy more broadly has found that when people have greater sense of control over their data, they are likely to be willing to share more data (e.g. [11, 68, 75]), even if this control is merely an artifact of transparency and not of actual usage of the data. Therefore, in this study, we build upon this prior work specifically in the health domain: we explore the role of messaging related to privacy control, and the transparency of the data being collected, on the likelihood to adopt a COVID-19 app, and how these factors intersect with how the appeal of the app is framed, and the socio-demographics of the adopter.

3 METHODS

To answer our research questions, we conducted a randomized, controlled field experiment using Google Ads. Here, we review our experimental design, data collection, ethical considerations, analysis approach, and the limitations of our work.

3.1 Experimental Design

Upon the public release of the CovidDefense app, we ran 14 separate Google display ad campaigns from February 1 to 26, 2021. In collaboration with the state of Louisiana, these were the only Google Display ads run for CovidDefense during that time. Each campaign was targeted, via IP address, at people who reside in Louisiana.⁷ All campaigns used the same settings, ad destination, and ad image from the state of Louisiana's CovidDefense marketing materials, and all campaigns were run concurrently. The 14 ads varied only in their text data in alignment with the 14 conditions summarized in Figure 1. Two examples of how an ad was presented to a user on a computer through Google Ads are depicted in Figure 2.

To evaluate existing policy guidance and findings from prior work, as summarized in the introduction, we chose to explore the factors of appeals, data transparency, and privacy transparency on contact tracing app adoption. Prior research finds additional adoption decision factors, such as who is operating the app/data infrastructure and whether the app is centralized or decentralized. We did not evaluate these additional factors as we were working with one app provider and a single app.

The image used in the ads were provided by the State of Louisiana's graphics design team. We used the same image on all of the ads. The layout of the ads is dictated by Google and thus the only items we could control were the ad text and ad image. The text of the 14 ads was chosen in the following manner. One of two appeals – individual ("Get notified of COVID exposure") or collective ("Reduce COVID infections") – appeared at the beginning of the ad text. These phrases were limited to 30 characters by the Google Ads platform. We selected these phrases based on a pilot test in collaboration with the state of Louisiana, in which a market research firm surveyed approximately 800 respondents to identify the best message phrasings that were most appealing on a variety of criteria. This allowed us to adopt already-successful messages and investigate, through this randomized study how the type of appeal,

as well as privacy and data transparency messaging, influenced app adoption.

Following the appeal, either one of three privacy-transparency statements was added, or there was no privacy statement. The privacy statements either was broadly stated ("...without harming your privacy"), or had a technical ("App data stays on your device.") or non-technical ("You control the data you share.") statement of control over privacy. Finally, a privacy statement could also have been paired with a data collection statement ("The app uses information about who you have been near."). The entire text of all 14 ads can be found in Appendix A.

3.2 Data Collection

We observe a total of 7,010,271 impressions on our ads. Google Ads does not allow for a user to limit impressions on campaigns, so we manually monitored campaign performance and aimed to stop each campaign at $500,000 \pm 35,000$ impressions, with an average of 500,733.6 impressions per campaign. In total, we observe 28,026 clicks on our 14 campaigns. The average Click Through Rate (CTR – the outcome of interest – the proportion between number of clicks and number of impressions) of the 14 campaigns was 0.398%, with standard deviation of 0.100%.

Along with the number of clicks and number of impressions, we also observe measures of demographics (age, gender and community density: urban vs. rural). Demographics are provided through Google Ads metadata. Age and gender are inferred by the Google Ads platform through past browsing behavior; the accuracy of these inferences has been validated against gold-standard social scientific probabilistic survey panels and other self-report data sources [46]. To label participants' community density we map participant counties (called Parishes in Louisiana), which are determined by Google Ads based on IP address, using the Census mapping to community density. 97.8% of the impressions (6,858,820) had an associated Parish while 55.9% of the impressions (3,920,232) had both age and gender labels.

3.3 Analysis

Our main analysis examines differences in click through rates (CTRs) for different ads based on their messaging text. For statements about statistical significance, we report $\alpha = 0.05$.

For RQ1, we perform an analysis with a two-sided two proportion z-test on the CTRs of collective-good and individual-good ads. This analysis uses all the impression data ($n=7,010,271$).

For RQ2-3, we run the regressions defined in Tables 2-5, using all the impression data ($n=7,010,271$). For RQ4, regarding demographics and geographics, we run the regressions defined in Tables 6-10. For privacy reasons, Google Ads separates demographic and location data and thus we cannot analyze age, gender, and community density with the entire dataset. As such, the regression models for RQ4 only analyze those impressions which have the relevant demographic ($n=3,920,232$) or geographic ($n=6,858,820$) information.

We report statistics as odds ratios for each regression. The odds ratio compares the ratio of odds for a baseline event to the odds for the contrasting event. Additionally, since the geographic and demographic data are a subset of our entire dataset (97.8% and 55.9% respectively), it is natural to be concerned that analyzing these

⁷Prior work has validated the accuracy of this state-level targeting [6].




Figure 1: Experimental design for the 14 messages shown in the field study.




Figure 2: Two examples of how Ad #1 is displayed on a computer.

subsets may lead to different conclusions. However, we analyzed the results of RQ2 through RQ4 with just the dataset subsets and found that the results are robust to such modeling specifications, defined by overlapping confidence intervals on the odds ratios for the same regressions with the different subsets. These results can be found in Tables 2-10 by comparing the "All data" column to the "Just Demographic" and "Just Geographic" columns.

Data Archival. Data and analysis scripts for this experiment are available at <https://doi.org/10.7910/DVN/MLUR6D>.

3.4 Ethics

Our study was approved by our institution's ethics review board and exempted from review by the Louisiana Department of Health IRB. All data collection occurred within the publicly available Google Ads platform and was explicitly approved by Google Ads. We only had access to information about users in the manner in which Google provided them. The data were presented to us in an aggregate manner, which preserved the users privacy in accordance

with Google Ads privacy policy.⁸ This aggregate data included geographic location of an impression with resolution to the Parish level. Google also provided some demographic information which are either user-supplied or inferred by browsing habits.

Further, we were careful to consider the potential real world effects of our ads. To that end, we were aware that some combinations of messaging could act as a deterrent to future adoption of these technologies. Since our research takes the view that adoption of contact tracing apps should be encouraged, we chose to not include combinations of messages that might work against that goal – that is why in our experimental design we only showed a data-transparency statement in combination with a privacy-transparency statement.

3.5 Limitations

Our results rely on a single study, from a single state (Louisiana) and with demographic data that rely on Google's ability to accurately classify gender, region, and age. The largest limitation of our

⁸<https://safety.google/privacy/ads-and-data/>.

study design is that we only capture clicks on ads instead of full downloads and app use. Our study was designed this way with a privacy focus and there is no way for us to reasonably measure conversion rate from click to download. However future studies could examine app adoption and use in conjunction with pro-social messaging questions. Additionally, while we chose the language used in our ad messages carefully, other forms of appeals or transparency statements could have been used. We encourage future studies to further investigate the impact of collective- vs. individual-goods appeals, as well as privacy and data transparency, on encouraging pro-social digital health behavior.

Further, we are limited in our demographic analysis by what the Google Ads platform provides us. While the demographic inferences made by Google Ads have been found to be very accurate [46], not all individuals in our dataset have inferred demographics and we are limited in our demographic analyses to the three demographics offered by Google Ads. We encourage further research on this topic with additional demographic data collected from consenting individuals. We do note that our sample is from a population of technology adopters who use the internet on mobile, desktop, or tablet platforms. Thus, our results should be interpreted as indicative of that population, however this is not necessarily a limitation as the CovidDefense and other contact tracing apps require at least as much technology adoption as we see in our sample.

Finally, at the time of data collection (February 1 to 26) Louisiana, and the whole United States, was seeing a decline in the number of reported COVID-19 cases, hospitalizations, and deaths, while simultaneously seeing vaccinations just starting to become available to the broader public. Since the situation was constant across all ads presented, we can conclude that the responses we observe measure human behavior to the messages presented in the ads conditioned on the state of the pandemic and media environment at the time.

4 RESULTS

Our experimental factors – appeal, privacy transparency, and data transparency – all significantly relate to CTR.

4.1 RQ1: Collective-good appeals are superior

Advertisements mentioning individual-good perform significantly worse ($\chi^2 = 598.54$; $df=1$, $p < 0.001$) than those that use collective-good appeals (CTR of 0.341% vs. 0.458%) in ads containing only an appeal, see Figure 3a. This is still the case even when controlling for other experimental factors (O.R.=0.745; Appendix Table 2), and the interactions between them (O.R.=0.880; $p < 0.001$; Appendix Table 3), demographics (O.R.=0.747; $p < 0.001$; Appendix Table 6), and community density (O.R.=0.744; $p < 0.001$; Appendix Table 6).

4.2 RQ2 and RQ3: Effect of transparency statements depends on appeal

The type of appeal (collective or individual) moderates the impact of the transparency statements. We conclude this by observing significant interactions between the appeal and the transparency statements in logistic regression models (Appendix Table 3). Subsequently, performing a regression on each appeal individually, we report the odds ratios and errors for the transparency statements for each appeal in Figure 3b.

Messages with a collective-good appeal have a higher CTR when they have additional privacy-transparency statements – all three of the statements have O.R.s that range from 1.106 to 1.203 (Appendix Table 4), all with $p < 0.001$ – but result in lower CTR responses when paired with a data-transparency statement, i.e., when ads explicitly mentions what data is being collected (O.R. = 0.911; $p < 0.001$; Appendix Table 4). On the other hand, messages with an individual-good appeal have a lower CTR when paired with a technical privacy-transparency statement (O.R. = 0.619; $p < 0.001$; Appendix Table 5). However, the broad and non-technical control privacy-transparency statements cannot be deemed to affect the CTR of the same individual-good appeal ($p = 0.070$ and $p = 0.396$, respectively). Moreover, contrary to the negative effect of data transparency on CTR in the collective-good appeal condition, when data collection is made transparent in the individual-good appeal condition, we observe a higher CTR than in messages without such data transparency (O.R.=1.08; $p < 0.001$; Appendix Table 5). We observe that, in a single regression model containing interactions between the transparency statements and the appeals (Appendix Table 3), a data-transparency statement reduces the difference in CTR between messages with collective- vs. individual-good appeals, while inclusion of a privacy statement increases the difference in CTR between messages with the two different appeals.

Table 1: Modeling the statement differences by Appeal and Gender

	Dependent variable:			
	Clicks			
	Collective-Good		Individual-Good	
	Male	Female	Male	Female
	(1)	(2)	(3)	(4)
Privacy.Broad	1.160 (1.049, 1.283) $p = 0.004^{***}$	0.993 (0.906, 1.087) $p = 0.877$	0.796 (0.707, 0.896) $p = 0.0002^{**}$	1.198 (1.076, 1.335) $p = 0.002^{**}$
NonTech.Control	1.110 (1.009, 1.222) $p = 0.032^*$	1.083 (0.986, 1.188) $p = 0.095$	1.018 (0.899, 1.152) $p = 0.784$	1.333 (1.185, 1.498) $p = 0.00001^{**}$
Technical.Control	1.236 (1.118, 1.366) $p = 0.00004^{**}$	0.980 (0.897, 1.071) $p = 0.660$	0.460 (0.402, 0.526) $p < 0.001^{**}$	0.827 (0.738, 0.927) $p = 0.002^{**}$
Data.Transparency	0.849 (0.795, 0.907) $p = 0.00001^{**}$	0.961 (0.909, 1.015) $p = 0.157$	1.275 (1.178, 1.379) $p < 0.001^{**}$	1.058 (0.987, 1.134) $p = 0.113$
Constant	0.004 (0.004, 0.005) $p < 0.001^{**}$	0.005 (0.005, 0.006) $p < 0.001^{**}$	0.003 (0.003, 0.004) $p < 0.001^{**}$	0.004 (0.004, 0.004) $p < 0.001^{**}$
Observations	916,470	1,111,417	1,026,261	866,084
Log Likelihood	-27,378.620	-36,413.650	-20,733.730	-24,634.080
Akaike Inf. Crit.	54,767.230	72,837.310	41,477.460	49,278.150

Note:

* $p < 0.05$; ** $p < 0.01$

4.3 RQ4: Demographic and geographic influences

Next, we consider demographic differences in responses to our CovidDefense advertisements. Thus far it has been debated, on the




Figure 3: Main results from regression models. Error bars of 95% confidence are given. (a) Overall CTRs for individual- and collective-good appeals. (b) Odds ratio for each transparency statement modeled twice for ads presented with each of the two appeals. The individual-good appeal (in red) is above the collective-good appeal (in blue). Regression tables are Appendix Tables 4 and 5. (c) Overall odds ratio for each demographic and geographic variable. Regression table is Appendix Table 6.




Figure 4: Click through rates (CTR) for each age group are reported with each curve representing an appeal and a gender. Collective-good appeals to females have no difference across the ages, whereas males see a significant drop in CTR in the middle ages. Individual-good appeals for females are non-decreasing across ages and generally flat for males. Regression tables are Appendix Table 7 and 8. Error bars are 95% confidence Clopper-Pearson intervals.

basis of self-report evidence [8, 13, 73], whether men are less likely to adopt pro-social behaviors such as mask wearing. In this work we offer field evidence of a gender difference in receptiveness to COVID-19 pro-social messaging and behavior: men are significantly less likely to click on ads for CovidDefense (O.R. men = 0.794; $p < 0.001$; Appendix Table 6). We find that this effect varies in size based on the appeal shown in the ad. Both men and women consistently prefer collective-goods ads and men are less likely to click on both collective- and individual-goods ads. However, the gap in CTR between men and women is significantly larger for individual-goods ads: males click 23% less often than females when shown individual-good ads compared to 10% less for collective-good ads (with $p < 0.001$ for both; see Appendix Tables 7 and 8).

Overall, we find that users between 34-64 are significantly less likely to click on the advertisements than those between 18 and 24 (O.R.s range from 0.874 - 0.951 with $p < 0.05$; see Figure 3c and Appendix Table 6). On the other hand, those over 65 – who are also at high risk for developing and dying from COVID-19 [19] – are significantly more likely to click than those 18-24 (O.R. = 1.13; $p < 0.001$; Appendix Table 6). However, these effects are moderated both by the appeal of the ad shown and the gender of the ad viewer. Specifically, CTRs do not vary among different age groups of women when shown a collective-good ad (O.R.s range from 0.937-1.08 with $p > 0.1$; Appendix Table 7), likely because of the strength of this appeal for women, whom prior research shows are more relational – focused on the collective good – than men [34]. On the other

hand, when shown an individual-good appeal, women 45 to 65+ are significantly more likely to click than younger women (each age group above 45 sees an increase in CTR over the previous; O.R.s range from 1.34-1.58 with $p < 0.001$; Appendix Table 8). There is no difference between CTRs for ages 18-44 (O.R.s = 1.02 with $p > 0.7$; Appendix Table 8). We hypothesize that, aligned with [18], older women are cognizant of their higher COVID-19 risk and thus are willing to click even on an ad with an appeal that is less preferable and does not align with their broad tendency toward relationally-guided behavior.

Among men who were shown a collective-good appeal, young men (18-24 years old) and older men (65+ years old) are equally likely to click when shown a collective-good ad (O.R.=1.02; $p = 0.728$; Appendix Table 7), and are significantly more likely to click than middle-aged men (O.R.s range from 0.597-0.849; $p < 0.001$; Appendix Table 7). Considering individual-good ads, we observe a similar pattern, with no significant differences in the likelihood of clicking among men aged 18-44 and those over 65 (O.R.s range from 0.915-1.07; $p > 0.1$; Appendix Table 8), while men aged 45-64 are significantly less likely to click on the same ads (O.R.s range from 0.760-0.785; $p < 0.001$; Appendix Table 8). We find that men, especially those who are middle-aged and when presented with individual-good appeal, are far less likely than women to click to install CovidDefense.

Prior literature finds that men have lower perceptions of their COVID-19 risk [18]. We hypothesize that the large gap in CTR between men and women, which is especially pronounced when presented with an individual-good appeal, is driven from the gender-based risk tolerance differences documented in the literature. When presented with an individual-good appeal that primes the viewer to especially focus on their own risk, the gender differences are even more pronounced.

Beyond the moderating effects of age and gender on the appeal used to advertise CovidDefense, we also find gender differences in the effect of our privacy- and data-transparency statements (see Table 1). While the effects shown in Figure 3b are consistent across ages, we find that the overall effects of the privacy- and data-transparency statements are primarily driven by men's response to these statements. When shown a collective-good appeal, men are more likely to click if presented with privacy controls of any sort, but are less likely to click if there is an explicit data-transparency statement. Women, on the other hand, show different responses to the transparency statements. Specifically, when combined with a collective-good appeal, both privacy- and data-transparency statements have no impact on women's likelihood to click.

Further, while both men and women are less likely to click on individual-good ads that include a technical privacy control, this reduced likelihood to click is larger among men (CI for O.R. for men: (0.402, 0.526); CI for O.R. for women: (0.738, 0.927); Table 1). In addition, while men also respond negatively, or have no significant response, to the non-technical privacy-transparency statements, women exhibit higher CTRs when these statements are paired with the individual-good appeal. Additionally, women are unaffected by the inclusion of a data-transparency statement compared to men who are more likely to click ads with an individual-good appeal when a data-transparency statement is included.

Finally, we examine the impact of geography on willingness to click on ads for CovidDefense; see Appendix Table 10. In contrast to the existing body of literature around urban-rural differences in COVID-19 behavior, which finds that rural residents are less concerned about COVID-19 and less likely to adopt pro-social COVID-19 behaviors [12, 15, 27, 29], we find that Louisiana residents in rural communities are significantly more likely to click on any of the proposed ads (O.R. = 1.15; $p < 0.001$; Appendix Table 6). This finding is robust across both appeals and all transparency statements with the appeal preference also being robust between geographies; we observe that rural residents also find that collective-good statements as preferable to individual-good appeals.

5 DISCUSSION

The results of our work offer implications for thinking about transparent messaging in the promotion of privacy-sensitive health technologies. Given that the U.S. is a highly individualistic country, it is perhaps surprising that collective-good appeals were more effective than individual-good appeals in encouraging people in Louisiana to click to adopt the CovidDefense app. Prior work finds that Louisiana is the most collectivist U.S. state [1]. Collectivists tend to engage in pro-social behavior that benefits the in-group, rather than pro-social behavior that they perceive as benefiting themselves individually [4, 35]. This effect may explain why residents of Louisiana respond best to societally-oriented benefits. On the other hand, it may simply be the case that people understand that the primary benefit of a COVID-19 exposure notification app is indeed collective and that people respond best to messages that are honest and transparent about the true benefit of the app.

Our findings for transparency regarding privacy and data collection are, however, more nuanced. Transparency about individual data collection improves the efficacy of messages that are already individually focused – those with individual-good appeals – while the same transparency statement applied in a collectivist setting appears to conflict with people's sense of collective purpose. Relatedly, transparency regarding privacy and how individual data is protected in a collective setting may reduce concerns about personal privacy risk in a communal context, which prior work finds may be especially elevated, while the same transparency may be ineffective or even detrimental when placed in the context of individualistic privacy-benefit trade-offs [67].

Complicating our transparency findings, we observe gender effects in response to both the appeals and the privacy/data transparency statements. Overall, we find that men, especially those who are middle-aged and when presented with individual-good appeal, are far less likely than women to click to install CovidDefense. Prior literature finds that men have lower perceptions of their COVID-19 risk [18]. We hypothesize that the large gap in CTR between men and women, which is especially pronounced when presented with an individual-good appeal, is driven from the gender-based risk tolerance differences documented in the literature. When presented with an individual-good appeal that primes the viewer to especially focus on their own risk, the gender differences are becoming more pronounced. Differences between men and women also exist for the privacy- and data-transparency statements. The overall transparency effects described above are primarily driven

by men's response to these statements. Women are unaffected by the inclusion of privacy and data-transparency statements in the more-effective collective-good ads. We hypothesize that this is the case because women are more relational – focused on the collective good – than men [34] and thus, similar to the lack of age effect observed for women when presented with collective-good appeals, we hypothesize that there is such strong alignment between women's tendency toward relational choices and the collective-good appeal that other factors (e.g., age-specific risk perceptions, transparency statements) have no significant effect.

When presented with an individual-good appeal, both men and women are less likely to click when a technical privacy-transparency statement is included; the size of this effect is significantly larger for men. However, women are *more* likely to click when a less-technical privacy statement is paired with an individual-good appeal; while men are more likely to click – and women are less likely to – when a data-transparency statement is added. Taken together, these findings – that women are affected positively by less technical statements of privacy while men are positively affected by technical and data related privacy statements – align with prior work finding that men and women focus on different privacy controls: men have been found to focus more on technical privacy controls, while women are more focused on privacy sentiment and non-technical controls [26, 39, 45, 57].

In sum, our results suggest that a one-size-fits-all approach to advertising digital health applications may explain existing failures of such applications, including COVID-19 exposure notification apps, to achieve their potential [42, 78]. Future digital health applications must carefully consider how to frame the benefits of adoption and how to balance an ethical duty to privacy and data transparency with the ways in which our findings reveal this transparency alters the privacy-benefit calculus of different groups of potential users.

6 CONCLUSION

We present the first field study of COVID-19 app adoption. In a large-scale randomized field study, we find that residents of Louisiana are more likely to click on ads for exposure notification apps if the ad included a collective-good appeal. This effect was moderated (especially for men) by transparency regarding the individual data being collected and privacy protections offered for that data, likely due to the conflict between the sense of collective purpose and the associated cost for individual privacy. Moreover, we find age and gender differences in the likelihood to click the ads, fitting with past literature on varying risks of COVID-19 across ages and gender differences in perceptions of COVID-19 risk. These differences included lower probability to click on ads with technical privacy controls for women and higher likelihood for older people to click on ads due to the higher risk for COVID-19 associated with older people. We also find that the gender differences were made larger when the ads were individually-focused (with individual-good appeal), suggesting that the priming for individualism enhances gender differences. These findings may aid companies and policy makers when promoting digital tools to improve public health, especially those tools that have implications for privacy.

REFERENCES

- [1] Jüri Allik and Anu Realo. 2004. Individualism-collectivism and social capital. *Journal of cross-cultural psychology* 35, 1 (2004), 29–49.
- [2] Susan Athey, Christian Catalini, and Catherine Tucker. 2017. *The digital privacy paradox: Small money, small costs, small talk*. Technical Report. National Bureau of Economic Research.
- [3] Ali Balapour, Iris Reychav, Rajiv Sabherwal, and Joseph Azuri. 2019. Mobile technology identity and self-efficacy: Implications for the adoption of clinically supported mobile health apps. *International Journal of Information Management* 49 (2019), 58–68.
- [4] Delia Baldassarri and Maria Abascal. 2020. Diversity and prosocial behavior. *Science* 369, 6508 (2020), 1183–1187.
- [5] Matthias Baldauf, Peter Fröhlich, and Rainer Endl. 2020. Trust Me, I'm a Doctor—User Perceptions of AI-Driven Apps for Mobile Health Diagnosis. In *19th International Conference on Mobile and Ubiquitous Multimedia*. 167–178.
- [6] Jack Bandy and Brent Hecht. 2021. Errors in Geotargeted Display Advertising: Good News for Local Journalism? *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW (2021).
- [7] Sachin Banker and Joowon Park. 2020. Evaluating prosocial COVID-19 messaging frames: Evidence from a field study on Facebook. *Judgment and Decision Making* 15, 6 (2020), 1037–1043.
- [8] Sarah J Barber and Hyunji Kim. 2021. COVID-19 worries and behavior changes in older and younger men and women. *The Journals of Gerontology: Series B* 76, 2 (2021), e17–e23.
- [9] Matt Begun, Anthony T Newall, Guy B Marks, and James G Wood. 2013. Contact tracing of tuberculosis: a systematic review of transmission modelling studies. *PLoS One* 8, 9 (2013), e72470.
- [10] Yoshua Bengio, Daphne Ippolito, Richard Janda, Max Jarvie, Benjamin Prud'homme, Jean-François Rousseau, Abhinav Sharma, and Yun William Yu. 2021. Inherent privacy limitations of decentralized contact tracing apps. *Journal of the American Medical Informatics Association* 28, 1 (2021), 193–195.
- [11] Laura Brandimarte, Alessandro Acquisti, and George Loewenstein. 2013. Misplaced confidences: Privacy and the control paradox. *Social psychological and personality science* 4, 3 (2013), 340–347.
- [12] Timothy Callaghan, Jennifer A Lueck, Kristin Lunz Trujillo, and Alva O Ferdinand. 2021. Rural and urban differences in COVID-19 prevention behaviors. *The Journal of Rural Health* (2021).
- [13] Dan Cassino and Yasemin Besen-Cassino. 2020. Of Masks and Men? Gender, Sex, and Protective Measures during COVID-19. *Politics & Gender* 16, 4 (2020), 1052–1062.
- [14] Eugene Y Chan and Najam U Saqib. 2021. Privacy concerns can explain unwillingness to download and use contact tracing apps when COVID-19 concerns are high. *Computers in Human Behavior* 119 (2021), 106718.
- [15] Xuewei Chen and Hongliang Chen. 2020. Differences in preventive behaviors of COVID-19 between urban and rural residents: lessons learned from a cross-sectional study in China. *International journal of environmental research and public health* 17, 12 (2020), 4437.
- [16] Mihail Cocosila and Norm Archer. 2010. Adoption of mobile ICT for health promotion: an empirical investigation. *Electronic Markets* 20, 3 (2010), 241–250.
- [17] Zhaoxia Deng, Ziyang Hong, Cong Ren, Wei Zhang, and Fei Xiang. 2018. What predicts patients' adoption intention toward mHealth services in China: empirical study. *JMIR mHealth and uHealth* 6, 8 (2018), e172.
- [18] Ying Fan, A Yeşim Orhun, and Dana Turjeman. 2020. *Heterogeneous actions, beliefs, constraints and risk tolerance during the COVID-19 pandemic*. Technical Report. National Bureau of Economic Research.
- [19] Centers for Disease Control, Prevention, et al. 2020. Older adult at greater risk of requiring hospitalization or dying if diagnosed with COVID-19.
- [20] Linux Foundation. [n.d.]. Linux Foundation Public Health Landscape. <https://landscape.lfph.io/>. (Accessed on 08/19/2021).
- [21] Jemima A. Frimpong and Stephane HELLERINGER. 2020. *Financial Incentives for Downloading COVID-19 Digital Contact Tracing Apps*. preprint. SocArXiv. <https://doi.org/10.31235/osf.io/9vp7x>
- [22] Thomas Fritz, Elaine M Huang, Gail C Murphy, and Thomas Zimmermann. 2014. Persuasive technology in the real world: a study of long-term use of activity sensing devices for fitness. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 487–496.
- [23] Sarah Geber and Thomas Friemel. 2021. A Typology-Based Approach to Tracing-App Adoption During the COVID-19 Pandemic: The Case of the SwissCovid App. *Journal of Quantitative Description: Digital Media* 1 (April 2021). <https://doi.org/10.51685/jqd.2021.007>
- [24] Gilie Gefen, Omer Ben-Porat, Moshe Tennenholtz, and Elad Yom-Tov. 2020. Privacy, altruism, and experience: Estimating the perceived value of Internet data for medical uses. In *Companion Proceedings of the Web Conference 2020*. 552–556.
- [25] Jeremy Ginsberg, Matthew H Mohebbi, Rajan S Patel, Lynnette Brammer, Mark S Smolinski, and Larry Brilliant. 2009. Detecting influenza epidemics using search engine query data. *Nature* 457, 7232 (2009), 1012–1014.

- [26] Hana Habib, Pardis Emami Naeini, Summer Devlin, Maggie Oates, Chelse Swoopes, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. 2018. User behaviors and attitudes under password expiration policies. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS) 2018*. 13–30.
- [27] Michael H Haischer, Rachel Beilfuss, Meggie Rose Hart, Lauren Opielinski, David Wucke, Gretchen Zirgaitis, Toni D Uhrich, and Sandra K Hunter. 2020. Who is wearing a mask? Gender-, age-, and location-related differences during the COVID-19 pandemic. *PloS one* 15, 10 (2020), e0240785.
- [28] Laszlo Horvath, Susan Banducci, and Oliver James. 2020. Citizens' Attitudes to Contact Tracing Apps. *Journal of Experimental Political Science* (Sept. 2020), 1–13. <https://doi.org/10.1017/XPS.2020.30>
- [29] Qian Huang, Sarah Jackson, Sahar Derakhshan, Logan Lee, Erika Pham, Amber Jackson, and Susan L Cutter. 2021. Urban-rural differences in COVID-19 exposures and outcomes in the South: A preliminary analysis of South Carolina. *PloS one* 16, 2 (2021), e0246548.
- [30] Maia L Jacobs, James Clawson, and Elizabeth D Mynatt. 2015. Comparing health information sharing preferences of cancer patients, doctors, and navigators. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 808–818.
- [31] Hannah Julienne, Ciarán Lavin, Cameron Belton, Martina Barjaková, Shane Timmons, and Peter D Lunn. 2020. Behavioural pre-testing of COVID Tracker, Ireland's contact-tracing app. (2020).
- [32] Gabriel Kapthuk, Daniel G. Goldstein, Eszter Hargittai, Jake Hofman, and Elissa M. Redmiles. 2020. How good is good enough for COVID19 apps? The influence of benefits, accuracy, and privacy on willingness to adopt. *arXiv:2005.04343 [cs]* (May 2020). <http://arxiv.org/abs/2005.04343> arXiv: 2005.04343.
- [33] Maria Karampela, Sofia Ouhbi, and Minna Isomursu. 2019. Connected health user willingness to share personal health data: questionnaire study. *Journal of medical Internet research* 21, 11 (2019), e14537.
- [34] Yoshihisa Kashima, Susumu Yamaguchi, Uichol Kim, Sang-Chin Choi, Michele J Gelfand, and Masaki Yuki. 1995. Culture, gender, and self: a perspective from individualism-collectivism research. *Journal of personality and social psychology* 69, 5 (1995), 925.
- [35] Markus Kemmelmeier, Edina E Jambor, and Joyce Letner. 2006. Individualism and good works: Cultural variation in giving and volunteering across the United States. *Journal of Cross-Cultural Psychology* 37, 3 (2006), 327–344.
- [36] Tami Kim, Kate Barasz, and Leslie K John. 2019. Why am I seeing this ad? The effect of ad transparency on ad effectiveness. *Journal of Consumer Research* 45, 5 (2019), 906–932.
- [37] Predrag Klasnja, Sunny Consolvo, Tanzeem Choudhury, Richard Beckwith, and Jeffrey Hightower. 2009. Exploring privacy concerns about personal sensing. In *International Conference on Pervasive Computing*. Springer, 176–183.
- [38] Lars Korn, Robert Böhm, Nicolas W Meier, and Cornelia Betsch. 2020. Vaccination as a social contract. *Proceedings of the National Academy of Sciences* 117, 26 (2020), 14890–14899.
- [39] Feng-Yang Kuo, Cathy S Lin, and Meng-Hsiang Hsu. 2007. Assessing gender differences in computer professionals' self-regulatory efficacy concerning information privacy practices. *Journal of business ethics* 73, 2 (2007), 145–160.
- [40] Betty Ladyzhets. 2021. We investigated whether digital contact tracing actually worked in the US. <https://www.technologyreview.com/2021/06/16/1026255/us-digital-contact-tracing-exposure-notification-analysis/>
- [41] Marc Langheinrich and Florian Schaub. 2018. Privacy in mobile and pervasive computing. *Synthesis Lectures on Mobile and Pervasive Computing* 10, 1 (2018), 1–139.
- [42] Dyani Lewis. 2020. Why many countries failed at COVID contact-tracing-but some got it right. *Nature* (2020), 384–387.
- [43] Tianshi Li, Camille Cobb, Jackie Yang, Sagar Baviskar, Yuvraj Agarwal, Beibei Li, Lujo Bauer, and Jason I Hong. 2021. What makes people install a COVID-19 contact-tracing app? Understanding the influence of app design and individual difference on contact-tracing app adoption intention. *Pervasive and Mobile Computing* (2021), 101439.
- [44] Steven Lockey, Martin R Edwards, Matthew J Hornsey, Nicole Gillespie, Saeed Akhlaghpour, and Shannon Colville. 2021. Profiling adopters (and non-adopters) of a contact tracing mobile application: insights from Australia. *International Journal of Medical Informatics* 149 (2021), 104414.
- [45] Arunesh Mathur, Jessica Vitak, Arvind Narayanan, and Marshini Chetty. 2018. Characterizing the use of browser-based blocking extensions to prevent online tracking. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS) 2018*. 103–116.
- [46] Paul McDonald, Matt Mohebbi, and Brett Slatkin. 2012. Comparing google consumer surveys to existing probability and non-probability based internet surveys. *Google White Paper* (2012).
- [47] Christina Munns and Subhajit Basu. 2017. *Privacy and healthcare data: Choice of Control to 'Choice' and 'Control'*. Routledge.
- [48] Simon Munzert, Peter Selb, Anita Gohdes, Lukas F Stoetzer, and Will Lowe. 2021. Tracking and promoting the usage of a COVID-19 contact tracing app. *Nature Human Behaviour* 5, 2 (2021), 247–255.
- [49] National Academies of Sciences, Engineering, and Medicine and others. 2020. Encouraging Adoption of Protective Behaviors to Mitigate the Spread of COVID-19: Strategies for Behavior Change.
- [50] Alfred Ng. 2021. Google Promised Its Contact Tracing App Was Completely Private - But It Wasn't. <https://themarkup.org/privacy/2021/04/27/google-promised-its-contact-tracing-app-was-completely-private-but-it-wasnt>. (Accessed on 08/19/2021).
- [51] Andreia Nunes, Teresa Limpo, and São Luís Castro. 2019. Acceptance of mobile health applications: examining key determinants and moderators. *Frontiers in psychology* 10 (2019), 2791.
- [52] Jonathan A Obar and Anne Oeldorf-Hirsch. 2020. The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society* 23, 1 (2020), 128–147.
- [53] Aarathi Prasad, Jacob Sorber, Timothy Stablein, Denise Anthony, and David Kotz. 2012. Understanding sharing preferences and behavior for mHealth devices. In *Proceedings of the 2012 ACM workshop on Privacy in the electronic society*. 117–128.
- [54] Davy Preuveneers and Wouter Joosen. 2016. Privacy-enabled remote health monitoring applications for resource constrained wearable devices. In *Proceedings of the 31st Annual ACM Symposium on Applied Computing*. 119–124.
- [55] Nathaniel Rabb, David Glick, Attiyya Houston, Jake Bowers, and David Yokum. 2021. No evidence that collective-good appeals best promote COVID-related health behaviors. *Proceedings of the National Academy of Sciences* 118, 14 (2021).
- [56] Ramesh Raskar, Greg Nadeau, John Werner, Rachel Barbar, Ashley Mehra, Gabriel Harp, Markus Leopoldseider, Bryan Wilson, Derrick Flakoll, Praneeth Vepakomma, et al. 2020. COVID-19 contact-tracing mobile apps: evaluation and assessment for decision makers. *arXiv preprint arXiv:2006.05812* (2020).
- [57] Elissa Redmiles. 2018. Net benefits: Digital inequities in social capital, privacy preservation, and digital parenting practices of US social media users. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 12.
- [58] Elissa M. Redmiles. 2020. User Concerns & Tradeoffs in Technology-facilitated COVID-19 Response. *Digital Government: Research and Practice* 2, 1 (Nov. 2020), 6:1–6:12. <https://doi.org/10.1145/3428093>
- [59] Florian Schaub, Rebecca Balebako, and Lorrie Faith Cranor. 2017. Designing effective privacy notices and controls. *IEEE Internet Computing* (2017).
- [60] John S Seberger and Sameer Patil. 2021. Us and Them (and It): Social Orientation, Privacy Concerns, and Expected Use of Pandemic-Tracking Apps in the United States. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [61] Katrina J Serrano, Mandi Yu, William T Riley, Vaishali Patel, Penelope Hughes, Kathryn Marchesini, and Audie A Atienza. 2016. Willingness to exchange health information via mobile devices: findings from a population-based survey. *The Annals of Family Medicine* 14, 1 (2016), 34–40.
- [62] Alessio Signorini, Alberto Maria Segre, and Philip M Polgreen. 2011. The use of Twitter to track levels of disease activity and public concern in the US during the influenza A H1N1 pandemic. *PloS one* 6, 5 (2011), e19467.
- [63] Lucy Simko, Jack Lucas Chang, Maggie Jiang, Ryan Calo, Franziska Roesner, and Tadayoshi Kohno. 2020. COVID-19 Contact Tracing and Privacy: A Longitudinal Study of Public Opinion. *arXiv:2012.01553 [cs.CY]*
- [64] Elizabeth Stowell, Mercedes C Lyson, Herman Saksono, René C Wurth, Holly Jimison, Misha Pavel, and Andrea G Parker. 2018. Designing and evaluating mHealth interventions for vulnerable populations: A systematic review. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [65] Yongqiang Sun, Nan Wang, Xitong Guo, and Zeyu Peng. 2013. Understanding the acceptance of mobile health services: a comparison and integration of alternative models. *Journal of electronic commerce research* 14, 2 (2013), 183.
- [66] Séverine Toussaert. 2021. Upping uptake of COVID contact tracing apps. *Nature Human Behaviour* 5, 2 (2021), 183–184.
- [67] Sabine Treppe, Leonard Reinecke, Nicole B Ellison, Oliver Quiring, Mike Z Yao, and Marc Ziegele. 2017. A cross-cultural perspective on the privacy calculus. *Social Media+ Society* 3, 1 (2017), 2056305116688035.
- [68] Catherine E Tucker. 2014. Social networks, personalized advertising, and privacy controls. *Journal of marketing research* 51, 5 (2014), 546–562.
- [69] Felix Velicia-Martin, Juan-Pedro Cabrera-Sanchez, Eloy Gil-Cordero, and Pedro R. Palos-Sanchez. 2021. Researching COVID-19 tracing app acceptance: incorporating theory from the technological acceptance model. *PeerJ Computer Science* 7 (Jan. 2021), e316. <https://doi.org/10.7717/peerj-cs.316>
- [70] Viswanath Venkatesh, Michael G Morris, Gordon B Davis, and Fred D Davis. 2003. User acceptance of information technology: Toward a unified view. *MIS quarterly* (2003), 425–478.
- [71] Michel Walrave, Cato Waeterloos, and Koen Ponnet. 2020. Adoption of a Contact Tracing App for Containing COVID-19: A Health Belief Model Approach. *JMIR Public Health and Surveillance* 6, 3 (Sept. 2020), e20572. <https://doi.org/10.2196/20572>
- [72] Mark Warner, Andreas Gutmann, M Angela Sasse, and Ann Blandford. 2018. Privacy unraveling around explicit HIV status disclosure fields in the online geosocial hookup app Grindr. *Proceedings of the ACM on human-computer interaction* 2, CSCW (2018), 1–22.

- [73] Alan White. 2020. Men and COVID-19: the aftermath. *Postgraduate Medicine* 132, sup4 (2020), 18–27.
- [74] Simon N. Williams, Christopher J. Armitage, Tova Tampe, and Kimberly Dienes. 2021. Public attitudes towards COVID-19 contact tracing apps: A UK-based focus group study. *Health Expectations* 24, 2 (April 2021), 377–385. <https://doi.org/10.1111/hex.13179>
- [75] Heng Xu, Hock-Hai Teo, Bernard CY Tan, and Ritu Agarwal. 2012. Research note—effects of individual self-protection, industry self-regulation, and government regulation on privacy concerns: a study of location-based services. *Information Systems Research* 23, 4 (2012), 1342–1363.
- [76] Baobao Zhang, Sarah Kreps, Nina McMurry, and R. Miles McCain. 2020. Americans' perceptions of privacy and surveillance in the COVID-19 pandemic. *PLOS ONE* 15, 12 (Dec. 2020), e0242652. <https://doi.org/10.1371/journal.pone.0242652>
- [77] Xiaofei Zhang, Xitong Guo, Kee-hung Lai, Feng Guo, and Chenlei Li. 2014. Understanding gender differences in m-health adoption: a modified theory of reasoned action model. *Telemedicine and e-Health* 20, 1 (2014), 39–46.
- [78] Bettina Maria Zimmermann, Amelia Fiske, Barbara Prainsack, Nora Hangel, Stuart McLennan, and Alena Buyx. 2021. Early perceptions of COVID-19 contact tracing apps in German-speaking countries: comparative mixed methods study. *Journal of medical Internet research* 23, 2 (2021), e25525.

APPENDIX

A ADVERTISEMENTS

The full text of the 14 ad campaigns are included here for clarity:

- (1) Get notified of COVID exposure.
- (2) Get notified of COVID exposure. CovidDefense uses information about who you have been near, without harming your privacy.
- (3) Get notified of COVID exposure. CovidDefense uses information about who you have been near. App data stays on your device.
- (4) Get notified of COVID exposure. CovidDefense uses information about who you have been near. You control the data you share.
- (5) Get notified of COVID exposure, without harming your privacy.
- (6) Get notified of COVID exposure. App data stays on your device.
- (7) Get notified of COVID exposure. You control the data you share.
- (8) Reduce COVID infections.
- (9) Reduce COVID infections. The app uses information about who you have been near, without harming your privacy.
- (10) Reduce COVID infections. The app uses information about who you have been near. App data stays on your device.
- (11) Reduce COVID infections. The app uses information about who you have been near. You control the data you share.
- (12) Reduce COVID infections, without harming your privacy.
- (13) Reduce COVID infections. App data stays on your device.
- (14) Reduce COVID infections. You control the data you share.

B DATASETS

The entire dataset and model regressions can be found here: <https://doi.org/10.7910/DVN/MLUR6D>.

There are two primary datasets: one which has all 7,010,271 impressions and demographic data, and another with just the impressions that have associated geographic information. The former includes columns for Google-estimated demographics like Age and Gender, with many impressions having values of "Unknown".

These two data tables for demographic and geographic impressions were represented by a row for each impression with columns for whether that impression resulted in a click; the age and gender or geography of the impression; as well as indicator variables for the presence or absence of ad information (appeals, privacy transparency – broad privacy reassurance, non-technical control, and technical control – and data transparency).

An associated R file is included which includes functions to reproduce each model and associated statistics.

C REGRESSIONS

We report the full regression tables for all claims made in the paper.

Table 2 includes overall effects for each experimental variable. These regressions are performed with all the data, just demographic data, and just geographic data.

Table 3 reports the interaction effect between each privacy and transparency statement with the two appeals. These regressions

are performed with all the data, just demographic data, and just geographic data.

Table 4 reports the regression for each statement for all **collective-good** ads. These regressions are performed with all the data, just demographic data, and just geographic data.

Table 5 reports the regression for each statement for all **individual-good** ads. These regressions are performed with all the data, just demographic data, and just geographic data.

Table 6 reports the regression for each experimental, demographic, and geographic variable.

Table 7 reports the regressions for age and gender differences for **collective-good** ads.

Table 8 reports the regressions for age and gender differences for **individual-good** ads.

Table 9 reports the regressions for the privacy and transparency statements, separated by appeal and gender combinations.

Table 10 reports the regressions for Urban/Rural interactions with the experimental variables.

Table 2: Modeling the five independent variables

	<i>Dependent variable:</i>		
	Clicks		
	All data	Just Demographic	Just Geographic
	(1)	(2)	(3)
Individual.Good	0.745 (0.727, 0.763) p < 0.001**	0.730 (0.707, 0.752) p < 0.001**	0.744 (0.726, 0.762) p < 0.001**
Privacy.Broad	1.032 (0.992, 1.074) p = 0.121	1.031 (0.980, 1.085) p = 0.243	1.034 (0.993, 1.076) p = 0.106
NonTech.Control	1.084 (1.042, 1.128) p = 0.0001**	1.112 (1.056, 1.172) p = 0.0001**	1.092 (1.049, 1.137) p = 0.00002**
Technical.Control	0.920 (0.883, 0.958) p = 0.0001**	0.890 (0.844, 0.938) p = 0.00002**	0.924 (0.886, 0.963) p = 0.0002**
Data.Transparency	0.981 (0.957, 1.007) p = 0.147	1.028 (0.995, 1.062) p = 0.100	0.977 (0.953, 1.003) p = 0.081
Constant	0.005 (0.004, 0.005) p < 0.001**	0.005 (0.005, 0.005) p < 0.001**	0.005 (0.004, 0.005) p < 0.001**
Observations	7,010,271	3,920,232	6,858,820
Log Likelihood	-182,372.500	-109,490.000	-178,407.300
Akaike Inf. Crit.	364,756.900	218,992.000	356,826.500

Note:

*p<0.05; **p<0.01

Table 3: Modeling the interaction of the appeal with the privacy and transparency statements

	<i>Dependent variable:</i>		
	Clicks		
	All data	Just Demographic	Just Geographic
	(1)	(2)	(3)
Individual.Good	0.880 (0.827, 0.937) p = 0.0001**	0.763 (0.703, 0.829) p < 0.001**	0.878 (0.824, 0.935) p = 0.0001**
Privacy.Broad	1.106 (1.048, 1.167) p = 0.0003**	1.073 (1.003, 1.147) p = 0.041*	1.110 (1.052, 1.172) p = 0.0002**
NonTech.Control	1.123 (1.064, 1.185) p = 0.00003**	1.096 (1.026, 1.172) p = 0.007**	1.131 (1.071, 1.195) p = 0.00001**
Technical.Control	1.203 (1.140, 1.269) p < 0.001**	1.095 (1.025, 1.169) p = 0.008**	1.209 (1.146, 1.277) p < 0.001**
Data.Transparency	0.911 (0.881, 0.942) p < 0.00001**	0.913 (0.875, 0.952) p = 0.00003**	0.904 (0.875, 0.935) p < 0.001**
Individual.Good:Data.Transparency	1.185 (1.126, 1.248) p < 0.001**	1.280 (1.197, 1.369) p < 0.001**	1.194 (1.134, 1.258) p < 0.001**
Individual.Good:Privacy.Broad	0.855 (0.789, 0.927) p = 0.0002**	0.910 (0.820, 1.009) p = 0.075	0.851 (0.785, 0.923) p = 0.0002**
Individual.Good:NonTech.Control	0.914 (0.843, 0.990) p = 0.028*	1.041 (0.934, 1.160) p = 0.469	0.914 (0.843, 0.991) p = 0.030*
Individual.Good:Technical.Control	0.515 (0.473, 0.559) p < 0.001**	0.581 (0.521, 0.648) p < 0.001**	0.514 (0.472, 0.559) p < 0.001**
Constant	0.004 (0.004, 0.004) p < 0.001**	0.005 (0.005, 0.005) p < 0.001**	0.004 (0.004, 0.004) p < 0.001**
Observations	7,010,271	3,920,232	6,858,820
Log Likelihood	-182,159.100	-109,357.700	-178,196.900
Akaike Inf. Crit.	364,338.100	218,735.300	356,413.700

Note:

*p<0.05; **p<0.01

Table 4: Modeling the privacy and transparency statements for Collective-Good ads

	<i>Dependent variable:</i>		
	Clicks		
	All data	Just Demographic	Just Geographic
	(1)	(2)	(3)
Privacy.Broad	1.106 (1.048, 1.167) p = 0.0003**	1.073 (1.003, 1.147) p = 0.041*	1.110 (1.052, 1.172) p = 0.0002**
NonTech.Control	1.123 (1.064, 1.185) p = 0.00003**	1.096 (1.026, 1.172) p = 0.007**	1.131 (1.071, 1.195) p = 0.00001**
Technical.Control	1.203 (1.140, 1.269) p < 0.001**	1.095 (1.025, 1.169) p = 0.008**	1.209 (1.146, 1.277) p < 0.001**
Data.Transparency	0.911 (0.881, 0.942) p < 0.00001**	0.913 (0.875, 0.952) p = 0.00003**	0.904 (0.875, 0.935) p < 0.001**
Constant	0.004 (0.004, 0.004) p < 0.001**	0.005 (0.005, 0.005) p < 0.001**	0.004 (0.004, 0.004) p < 0.001**
Observations	3,523,339	2,027,887	3,446,697
Log Likelihood	-102,945.400	-63,818.110	-100,767.200
Akaike Inf. Crit.	205,900.900	127,646.200	201,544.400
<i>Note:</i>			*p<0.05; **p<0.01

Table 5: Modeling the privacy and transparency statements for Individual-Good ads

	<i>Dependent variable:</i>		
	Clicks		
	All data (1)	Just Demographic (2)	Just Geographic (3)
Privacy.Broad	0.946 (0.891, 1.004) p = 0.070	0.976 (1.003, 1.147) p = 0.547	0.945 (0.890, 1.004) p = 0.068
NonTech.Control	1.026 (0.967, 1.089) p = 0.396	1.141 (1.026, 1.172) p = 0.003**	1.034 (0.974, 1.098) p = 0.274
Technical.Control	0.619 (0.581, 0.660) p < 0.001**	0.636 (1.025, 1.169) p < 0.001**	0.621 (0.582, 0.663) p < 0.001**
Data.Transparency	1.080 (1.038, 1.123) p = 0.0002**	1.169 (0.875, 0.952) p < 0.001**	1.080 (1.038, 1.124) p = 0.0002**
Constant	0.004 (0.004, 0.004) p < 0.001**	0.004 (0.005, 0.005) p < 0.001**	0.004 (0.004, 0.004) p < 0.001**
Observations	3,486,932	1,892,345	3,412,123
Log Likelihood	-79,213.640	-45,539.550	-77,429.660
Akaike Inf. Crit.	158,437.300	91,089.100	154,869.300

Note:

*p<0.05; **p<0.01

Table 6: Modeling demographics and geographics

	<i>Dependent variable:</i>	
	Clicks	
	Just Demographic (1)	Just Geographic (2)
Age25 - 34	0.951 (0.906, 0.998) p = 0.041*	
Age35 - 44	0.932 (0.886, 0.980) p = 0.006**	
Age45 - 54	0.874 (0.825, 0.925) p = 0.00001**	
Age55 - 64	0.909 (0.866, 0.954) p = 0.0001**	
Age65+	1.134 (1.077, 1.194) p = 0.00001**	
GenderMale	0.794 (0.769, 0.819) p < 0.001**	
DensityRural		1.146 (1.104, 1.189) p < 0.001**
Individual.Good	0.747 (0.724, 0.770) p < 0.001**	0.744 (0.726, 0.762) p < 0.001**
Privacy.Broad	1.014 (0.964, 1.068) p = 0.584	1.033 (0.993, 1.076) p = 0.110
NonTech.Control	1.115 (1.058, 1.175) p = 0.00005**	1.091 (1.048, 1.136) p = 0.00003**
Technical.Control	0.857 (0.813, 0.903) p < 0.001**	0.923 (0.886, 0.962) p = 0.0002**
Data.Transparency	1.034 (1.001, 1.068) p = 0.047*	0.978 (0.953, 1.003) p = 0.084
Constant	0.006 (0.005, 0.006) p < 0.001**	0.005 (0.004, 0.005) p < 0.001**
Observations	3,920,232	6,858,820
Log Likelihood	-109,314.800	-178,382.400
Akaike Inf. Crit.	218,653.500	356,778.700

Note:

*p<0.05; ** p<0.01

Table 7: Modeling the Age and Gender differences for Collective-Good ads

	<i>Dependent variable:</i>		
	Clicks		
	Collective-Good	Collective-Good Female	Collective-Good Male
	(1)	(2)	(3)
Age25 - 34	0.902 (0.849, 0.958) p = 0.001**	0.967 (0.886, 1.055) p = 0.455	0.849 (0.781, 0.923) p = 0.0002**
Age35 - 44	0.921 (0.864, 0.981) p = 0.012*	1.078 (0.985, 1.180) p = 0.103	0.791 (0.723, 0.866) p < 0.00001**
Age45 - 54	0.808 (0.751, 0.870) p < 0.00001**	0.952 (0.856, 1.058) p = 0.360	0.695 (0.627, 0.771) p < 0.001**
Age55 - 64	0.798 (0.750, 0.850) p < 0.001**	0.937 (0.865, 1.014) p = 0.106	0.597 (0.533, 0.669) p < 0.001**
Age65+	0.969 (0.906, 1.035) p = 0.344	1.016 (0.933, 1.105) p = 0.718	1.021 (0.911, 1.144) p = 0.728
GenderMale	0.887 (0.851, 0.924) p < 0.001**		
Constant	0.006 (0.006, 0.006) p < 0.001**	0.005 (0.005, 0.006) p < 0.001**	0.006 (0.005, 0.006) p < 0.001**
Observations	2,027,887	1,111,417	916,470
Log Likelihood	-63,777.320	-36,412.990	-27,332.210
Akaike Inf. Crit.	127,568.600	72,837.980	54,676.420

Note:

*p<0.05; **p<0.01

Table 8: Modeling the Age and Gender differences for Individual-Good ads

	<i>Dependent variable:</i>		
	Clicks		
	Individual-Good	Individual-Good Female	Individual-Good Male
	(1)	(2)	(3)
Age25 - 34	1.046 (0.964, 1.134) p = 0.280	1.020 (0.906, 1.149) p = 0.740	1.073 (0.959, 1.199) p = 0.218
Age35 - 44	0.962 (0.886, 1.045) p = 0.357	1.016 (0.903, 1.144) p = 0.788	0.915 (0.816, 1.027) p = 0.133
Age45 - 54	0.997 (0.911, 1.091) p = 0.947	1.337 (1.175, 1.522) p = 0.00002**	0.760 (0.670, 0.863) p = 0.00003**
Age55 - 64	1.074 (0.995, 1.160) p = 0.068	1.353 (1.218, 1.502) p < 0.00001**	0.795 (0.707, 0.894) p = 0.0002**
Age65+	1.331 (1.227, 1.443) p < 0.001**	1.580 (1.418, 1.760) p < 0.001**	1.065 (0.936, 1.211) p = 0.342
GenderMale	0.685 (0.653, 0.719) p < 0.001**		
Constant	0.004 (0.004, 0.004) p < 0.001**	0.004 (0.003, 0.004) p < 0.001**	0.003 (0.003, 0.004) p < 0.001**
Observations	1,892,345	866,084	1,026,261
Log Likelihood	-45,551.090	-24,641.040	-20,858.890
Akaike Inf. Crit.	91,116.180	49,294.090	41,729.780

Note:

*p<0.05; **p<0.01

Table 9: Modeling the statement differences by Appeal and Gender

	<i>Dependent variable:</i>			
	Clicks			
	Collective-Good Male	Collective-Good Female	Individual-Good Male	Individual-Good Female
	(1)	(2)	(3)	(4)
Privacy.Broad	1.160 (1.049, 1.283) p = 0.004**	0.993 (0.906, 1.087) p = 0.877	0.796 (0.707, 0.896) p = 0.0002**	1.198 (1.076, 1.335) p = 0.002**
NonTech.Control	1.110 (1.009, 1.222) p = 0.032*	1.083 (0.986, 1.188) p = 0.095	1.018 (0.899, 1.152) p = 0.784	1.333 (1.185, 1.498) p = 0.00001**
Technical.Control	1.236 (1.118, 1.366) p = 0.00004**	0.980 (0.897, 1.071) p = 0.660	0.460 (0.402, 0.526) p = 0.000**	0.827 (0.738, 0.927) p = 0.002**
Data.Transparency	0.849 (0.795, 0.907) p = 0.00001**	0.961 (0.909, 1.015) p = 0.157	1.275 (1.178, 1.379) p = 0.000**	1.058 (0.987, 1.134) p = 0.113
Constant	0.004 (0.004, 0.005) p = 0.000**	0.005 (0.005, 0.006) p = 0.000**	0.003 (0.003, 0.004) p = 0.000**	0.004 (0.004, 0.004) p = 0.000**
Observations	916,470	1,111,417	1,026,261	866,084
Log Likelihood	-27,378.620	-36,413.650	-20,733.730	-24,634.080
Akaike Inf. Crit.	54,767.230	72,837.310	41,477.460	49,278.150

Note:

*p<0.05; **p<0.01

Table 10: Modeling the statement differences with an interaction for Density

	<i>Dependent variable:</i>				
	Appeal	Privacy: Broad	Clicks		
			Non-Technical Control	Technical Control	Data Transparency
	(1)	(2)	(3)	(4)	(5)
DensityRural	1.114 (1.061, 1.170) p = 0.00002**	1.132 (1.082, 1.183) p < 0.00001**	1.164 (1.113, 1.217) p < 0.0001**	1.151 (1.103, 1.202) p < 0.0001**	1.158 (1.103, 1.216) p < 0.0001**
Individual.Good	0.738 (0.719, 0.757) p < 0.0001**				
DensityRural:Individual.Good	1.067 (0.990, 1.150) p = 0.090				
Privacy.Broad		1.025 (0.997, 1.054) p = 0.076			
DensityRural:Privacy.Broad		1.055 (0.973, 1.143) p = 0.197			
NonTech.Control			1.110 (1.080, 1.141) p < 0.0001**		
DensityRural:NonTech.Control			0.959 (0.884, 1.040) p = 0.314		
Technical.Control				0.875 (0.850, 0.901) p < 0.0001**	
DensityRural:Technical.Control				0.998 (0.917, 1.086) p = 0.959	
Data.Transparency					0.977 (0.952, 1.002) p = 0.075
DensityRural:Data.Transparency					0.983 (0.912, 1.060) p = 0.656
Constant	0.005 (0.004, 0.005) p < 0.0001**	0.004 (0.004, 0.004) p < 0.0001**	0.004 (0.004, 0.004) p < 0.0001**	0.004 (0.004, 0.004) p < 0.0001**	0.004 (0.004, 0.004) p < 0.0001**
Observations	6,858,820	6,858,820	6,858,820	6,858,820	6,858,820
Log Likelihood	-178,438.100	-178,732.500	-178,707.500	-178,688.200	-178,733.900
Akaike Inf. Crit.	356,884.200	357,473.000	357,423.000	357,384.500	357,475.800

Note:

*p<0.05; **p<0.01