

# Design(ing) Fictions for Collective Civic Reporting of Privacy Harms

YUXI WU, Georgia Institute of Technology, USA and Northeastern University, USA

WILLIAM AGNEW, Carnegie Mellon University, USA

W. KEITH EDWARDS, Georgia Institute of Technology, USA

SAUVIK DAS, Carnegie Mellon University, USA

Individually-experienced privacy harms are often difficult to demonstrate and quantify, which impedes efforts for their redress. Their effects often appear small and are inconsistently documented, and they only become more obvious when aggregated over time and across populations. Taking a design fiction approach, we explore the design requirements and cultural ideals of a government-run system that empowers people to collectively report on and make sense of experiences of privacy harm from online behavioral advertising. Through the use of fictional inquiry, story completion, and comicboarding methods, delivered in an online survey with 50 participants, we found that participants had detailed conceptions of the user experience of such a tool, but wanted assurance that their labor and personal data would not be exploited further by the government if they contributed evidence of harm. We extrapolate these design insights to government-supported complaint-reporting platforms in other domains, finding multiple common design gaps that might disincentivize people to report experiences of harm, be they privacy-related or otherwise.

CCS Concepts: • **Security and privacy** → *Human and societal aspects of security and privacy*; • **Human-centered computing** → *HCI design and evaluation methods*; **Collaborative and social computing design and evaluation methods**.

Additional Key Words and Phrases: civic technology, privacy harms, online behavioral advertising, collective action, design fiction

## ACM Reference Format:

Yuxi Wu, William Agnew, W. Keith Edwards, and Sauvik Das. 2025. Design(ing) Fictions for Collective Civic Reporting of Privacy Harms. *Proc. ACM Hum.-Comput. Interact.* 9, 2, Article CSCW063 (April 2025), 26 pages. <https://doi.org/10.1145/3710961>

## 1 Introduction

People are frustrated and concerned about online privacy, but feel helpless to effect change [3]. Prior work points to the potential of collective action systems to empower people to collaboratively demand change [58, 59]. However, prior work [63] has also found that experts discount these collaboratively-authored demands because, in legal and regulatory contexts, change and/or redress comes only after unambiguous demonstration of acute harm. In the United States, demonstration of acute harm can be difficult in the privacy context because privacy harms are difficult to quantify and because the more nefarious effects of these harms build up, over time, through repeated exposure [16, 62]. What's needed, then, is a way to chronicle evidence of privacy harms across a population over time. Today, there exists no easy way for the population-at-large to report on everyday privacy

---

Authors' Contact Information: [Yuxi Wu](mailto:yuxi@ccs.neu.edu), [yuxi@ccs.neu.edu](mailto:yuxi@ccs.neu.edu), Georgia Institute of Technology, Atlanta, Georgia, USA and Northeastern University, Boston, Massachusetts, USA; [William Agnew](mailto:wagnew@andrew.cmu.edu), [wagnew@andrew.cmu.edu](mailto:wagnew@andrew.cmu.edu), Carnegie Mellon University, Pittsburgh, Pennsylvania, USA; [W. Keith Edwards](mailto:keith@cc.gatech.edu), [keith@cc.gatech.edu](mailto:keith@cc.gatech.edu), Georgia Institute of Technology, Atlanta, Georgia, USA; [Sauvik Das](mailto:sauvik@cmu.edu), [sauvik@cmu.edu](mailto:sauvik@cmu.edu), Carnegie Mellon University, Pittsburgh, Pennsylvania, USA.



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2573-0142/2025/4-ARTCSCW063

<https://doi.org/10.1145/3710961>

harms. How might we design a system that empowers people to report on and make sense of these harms?

In recent years, there has been a rise in civic technology platforms from government agencies to support citizen complaint-reporting—e.g., in city 311 portals [29]—and gathering of collective public sentiment in other domains. As a broader example, the Consumer Financial Protection Bureau (CFPB) operates a complaints website where consumers can report problems with financial products and services and receive a response from the CFPB. Such platforms also exist in other federal-level institutions within the United States, such as the Federal Trade Commission (FTC), the Federal Communications Commission (FCC), and the Occupational Safety and Health Administration (OSHA). However, while these platforms receive and process millions of complaints, people continue to mistrust the government to adequately represent their interests in bringing privacy-violating corporations to justice [12] and remedying their harms more broadly [18].

One reason people might feel this way is that these existing systems primarily operate as intake forms, where the sole action people can take is submitting a complaint; this limited scope also fails to capture people’s nuanced expectations for seeking help from the government. Are there other opportunities for people to participate in a harm-reporting ecosystem? How might these possibilities change how we think about existing systems? We propose design fiction as a way of better understanding the collective needs and desires of people when using such systems. Through a mixture of fictional inquiry, story completion, and comicboarding methods deployed in an online survey, we explore the potentials of a fictional government-run reporting tool that empowers users to report on and make sense of privacy harms. To contextualize our design fictions, we focus on arguably the most pervasive end-user touchpoint for digital privacy harms: online behavioral advertising (OBA). We ask the following research questions:

- RQ1** What are design requirements for everyday end-users to participate in collective civic reporting of privacy harms from online behavioral advertising?
- RQ2** What cultural ideals do people have around redress for reported privacy harms from online behavioral advertising, and the parties that carry out this redress?

We find that participants had detailed conceptions of what the user interface and user experience of such a tool could look like, such as the ability to support multiple types of evidence of harm, and ticket numbers to follow-up on claims. They also wanted guarantees against privacy risks for contributing reports containing sensitive information, and had vocal expectations of speed, transparency, and accountability for such a tool. However, participants felt ambivalent about relying on volunteer labor to make sense of the data contributed through the tool: while some expressed the potential for pride in volunteering and a duty to help others, others worried that such a tool would simply be another way for their data and labor to be exploited by the government.

While our design fictions focused specifically on *privacy* harms from online behavioral advertising, participants also provided rich insights into their ideals for government reporting tools for *consumer* harms at large. We thus synthesized, from our findings, seven key design principles for supporting people’s trust in ***Government Tools for Civic Harm-reporting (GoTCHas)***: (1) visible, upfront benefits; (2) timely, useful feedback; (3) contestability; (4) error prevention measures; (5) integration into everyday life; (6) consideration of social influence; and (7) diversity and flexibility of commitment. We use these principles to evaluate a sample of five existing GoTCHAs in other domains, such as the CFPB’s complaint reporting site, finding that these existing GoTCHAs fail to offer both transparency into how they resolve people’s complaints, as well as concrete benefits for people to contribute. We also find that existing GoTCHAs lack consideration of social influence in people’s motivations to contribute evidence and complaints, offer very limited ways for people to contribute or participate, and, as standalone websites, are poorly integrated into everyday life.

These limitations can disincentivize people to contribute to these systems, further contributing to people's feelings of helplessness with respect to what the government can and will do about digital harms—especially as they relate to privacy, but also more broadly.

In summary, we make the following contributions in this paper:

- An adaptation of comicboarding and story completion methods to generate fictional futures where collective civic reporting of privacy harms is a reality;
- A rich understanding of people's cultural ideals surrounding OBA, and the government's capacity to meet those ideals;
- A set of seven design principles for building effective collective civic harm-reporting systems, generated from those user stories; and,
- A preliminary application of these principles to evaluating existing systems in other domains.

## 2 The Privacy Harms Landscape

In this section, we construct an initial motivation of our exploration of collective civic reporting as a promising avenue for redressing privacy harms. First, we discuss the documented gaps between collective user privacy preferences and the opinions of security and privacy experts. We then examine the nature of harms from privacy violations as a legal concept, and how getting legal recognition of privacy harms can be a way to bridge the user-expert gap. Finally, we argue for online behavioral advertising (OBA) as a uniquely interesting candidate domain for systems aimed at collective civic reporting of privacy harms.

### 2.1 Collective User Responses to Privacy Violations

There has been extensive documentation of users' reactions to a variety of mass-event institutional privacy violations, from discovering their information was a part of a data breach [24, 40], to learning about new regulations in the news, to perceiving output from recommendation algorithms as being too specific or creepy [53, 57]. A study of the aftermath of the 2017 Equifax data breach illustrated that while people were aware of the risks resulting from this breach, they tended not to take protective actions because they did not know enough about the breach or because it was cost-prohibitive to do so [67]. A recurrent theme, also directly reported in a recent Pew survey [3], is that despite their worry, fear or anger in response to these privacy violations, users tend not to take further action to protect themselves or react to the institution behind the event.

When people *do* make the effort to collectively respond, such as through online petitions, efforts have tended to fizzle out. One possible hypothesis for why this might occur is that petitions are often composed by one or only a few individuals and might not be fully representative of those who sign the petition. With fragmented ownership over the petition, as Shaw et al. [51] describe, people who sign it have little motivation to continue caring about the outcome of that petition beyond just signing it. Recent work has attempted to understand if scaffolding the petition process with opinion surveys and simple voting mechanisms could encourage users to take more of a stake in the process and elicit more representative sets of demands from users. For example, Wu et al. [63] found that it was fairly straightforward for users to reach consensus, and that users felt both collective empathy for one another and enthusiasm for sharing their concerns given this scaffolding.

### 2.2 Recognition of Privacy Harms

Despite this promising result, Wu et al. [63] also found that the demands that users generated and voted for—reparations for institutional privacy violations, modifications to the algorithms powering

targeted ads, and formal apologies from offending institutions—were summarily dismissed by security and privacy experts across industry, academic, and policy environments. Specifically, experts noted that it could be difficult to directly attribute harm and malintent to violating institutions.

This attribution problem is a core issue in the broader landscape of *privacy harms*, theorized by both Calo [10] and Citron and Solove [16], who assert that harms from privacy violations are currently inconsistently recognized by courts, and that certain non-financial and non-physical harms from privacy violations should be as cognizable as financial and physical ones. To this end, there have been some stirrings of legal recognition. For example, as early as 1980, the Federal Trade Commission (FTC) has recognized that small harms, in aggregate, can be sufficiently substantial if suffered by a large number of people<sup>1</sup>. Flashier violations, such as the Equifax data breach in 2017, have resulted in heavy fines from the FTC, in part because evidence of resultant harms on users is fairly easy to understand and define: financial losses for both institutions and consumers. However, this evidence is collected in an ad-hoc manner only once a data breach has come to light rather than as a sustained effort to support cases against *ongoing* violations and harms. We argue that designing formal, sustained systems to gather evidence of these privacy harms on a *collective* level, echoing [25], is key to lending further legitimacy to them in both legal settings and beyond.

Rakova et al. [46] have also proposed a social imaginary framework, Terms-we-serve-with (TwSw), for addressing broader *algorithmic* harms, specifically through a lens of *algorithmic reparation*. The authors specifically highlight “complaint and algorithmic harms reporting” as one of five key dimensions of this framework, which can not only help us address existing current issues of harm, but also *anticipate* them and prepare for them in the future. Our design fictions, which we will discuss in detail in later sections, directly build upon this dimension of TwSw by asking users to ideate on specific user experiences and interfaces for harm reporting and reflect on their societal implications.

### 2.3 Online Behavioral Advertising (OBA)

While privacy harms from data breaches might be clearly defined, as with the aforementioned FTC example, those entailed by OBA are less so, whether it be from feeling embarrassed because a targeted ad revealed intimate information about oneself to others, or chilling effects due to the specificity of ads. The mundanity of OBA makes it a unique candidate for exploring how to report privacy harms as they manifest in-situ.

**2.3.1 Harms of OBA.** Online targeted ads are highly effective at engaging users to click. However, people’s myriad reasons for disliking online targeted ads are extensively-reported: they find them creepy, privacy-invasive, and disruptive [4, 27, 54, 56, 57, 65, 66]. Past literature documenting the harms of ad targeting has primarily focused on targeting based on political interests, which can limit user exposure to diverse viewpoints [6, 7] (a phenomenon reinforced and exacerbated by the ad delivery mechanisms themselves [1]). However, as more recent work has pointed out, what an ad algorithm deems as “interests” can easily be someone’s vulnerability, e.g., harms related to sensitive health topics like weight-loss ads [23] or mental and physical health conditions [13]. More recently, Wu et al. [62] extended this work by examining and taxonomizing *generalized* harms from OBA. The authors found that there were four distinct categories of harms from OBA that users reported as most salient and personally-impactful: psychological distress, loss of autonomy, constriction of user behavior, and algorithmic marginalization and traumatization.

---

<sup>1</sup>FTC Policy Statement on Unfairness. December 17, 1980. <https://www.ftc.gov/legal-library/browse/ftc-policy-statement-unfairness>

2.3.2 *OBA as a Candidate for Evidence-Gathering and Harm-Reporting.* Literary scholar Rob Nixon [45] coined the term “slow violence” to describe things like the normalization of seemingly small harms, consisting of “calamities that are slow and long lasting, calamities that patiently dispense their devastation while remaining outside our flickering attention spans”. Gak et al. [23] adapted this concept from its original context—environmental degradation and climate change—to OBA: the ubiquity and embeddedness of invasive targeted ads in people’s experiences online numbs them to their harms over time [50], and exposes their privacy to a death by a thousand cuts. Despite this challenge, people are still motivated to take small actions against OBA in the form of preventative measures—they use ad-blockers, VPNs, and private browsing sessions, and even avoid having real-life conversations near their personal mobile devices to avoid being “listened to”. Gorski [26] and later Korff [37] both have argued that in cases where it’s difficult for plaintiffs to provide evidence of privacy harms due to the privacy violations themselves themselves being secret or obscured—such as in cases of national intelligence—people might instead provide evidence of a “diversion of time or resources”. Documenting such evasive actions and ubiquitous negative experiences alike can create a rich well of evidence of lived OBA harm.

While other types of privacy rights or harms might feel more straightforward to measure systematically—e.g., data access and deletion rights, dark patterns—it is precisely because OBA harms are so *un*-straightforward that our design exploration is necessary: is OBA, so varied in its presence and impact, even viable for harm-reporting? The potential evidence that we described above suggests it may be.

### 3 Design Context and Approach

While OBA harms might be amenable to a structured system for civic reporting, the design space for such a system is still broad and warrants additional exploratory work. In this section, we describe a rich landscape of collective civic harm-reporting in other domains, and why such existing systems might still be inadequate. We then propose a mixture of design fiction methodologies to address the gaps of this context. Finally, we detail how we developed our design fictions.

#### 3.1 Design Context: GoTCHas

We situate our work in a growing landscape of government-based civic reporting and evidence-gathering platforms in other domains within the United States of America. We term this class of systems as *Government Tools for Civic Harm-reporting (GoTCHas)*.

At the highest level, government agencies in the United States often solicit comments from the public on matters of legislation or to understand how to prioritize their resources, through the cross-agency federal website Regulations.gov. However, to the average person, this website is virtually unknown, and even if an individual has heard of it, it can be overwhelming to navigate: the home page features a catch-all search bar with few affordances for how to find a relevant topic to comment on, and consequently, thousands of search results. And, if a person has somehow found a relevant topic, there is no guidance for how to write a public comment, what information to include, or the format of the submission. As such, these public comments—which might not even be relevant to reporting harm at all—are primarily authored by policy professionals and experts, either as representatives of large corporations or from the federal government itself.

Two existing platforms for *non-expert* people to report Internet harms and evidence of those harms come from the Federal Trade Commission (FTC), the primary regulator of privacy and data security within the United States. The FTC operates ReportFraud.ftc.gov and IdentityTheft.gov, for people to report “fraud, scams, and bad business practices” and “report identity theft and get a recovery plan”, respectively. On both websites, consumers are guided through a questionnaire about the type of complaint they are submitting, and then asked to include specific documents to

support their complaint. On ReportFraud.ftc.gov, the FTC claims it will share fraud reports with law enforcement agencies and provide helpful tips for consumers to protect themselves from fraud in the future. On IdentityTheft.gov, there are tutorials for consumers to create plans and checklists for the recovery process, including in cases of data breaches. The FTC also publishes aggregated infographics of trends in fraud reporting, as well as whether consumers got money back from the resolution of those fraud cases.

Another well-established example of a civic reporting tool comes from the Consumer Financial Protection Bureau (CFPB), which operates a website for consumers to submit complaints about financial services and products, such as bank accounts, credit cards, mortgages, and other types of loans. Similar to the FTC, consumers on the CFPB complaints website fill out a guided form and can attach specific types of supporting evidence for their complaints. The CFPB claims that it will submit the complaint to the offending company on behalf of the consumer, or to another appropriate federal agency for a response usually within 15 days, but up to 60 days. On the opposite end, the CFPB publishes a publicly-accessible, anonymized database of these complaints that consumers submit, with an interactive visualization dashboard and API to download the data. The CFPB database goes one step further than the FTC's reports, however, and allows consumers to read and download individual complaints that others have submitted, including their anonymized personal accounts and associated evidence.

Examples from other domains include the Federal Communications Commission (FCC)'s Consumer Inquiries and Complaint Center, which, in a similar manner, accepts consumer complaints about telecommunications services like internet service and television broadcasts. The FCC's website also provides a form for privacy complaints, but these are limited to those related to privacy concerns about internet or telephone service providers. The Occupational Safety and Health Administration (OSHA) also operates a similar complaint website for people to report workplace health and safety issues, with a publicly accessible database. At the local level, municipal governments often operate 311 data portals where residents can submit complaints about non-emergency incidents, such as noise complaints, road blockages, streetlight outages, etc.; these portals often have also have publicly accessible, downloadable databases of these complaints similar to the CFPB.

While these existing GoTCHAs can all claim receipt of (and, to an extent, addressing) millions of complaints, the percentage of Americans who felt they could trust the government to "do the right thing" has not surpassed 30% since 2007 [12]. The context of privacy offers a similarly bleak outlook: only about 3 out of 10 Americans feel that the government will hold the CEOs of social media companies accountable if they misuse or compromise users' data, and 6 out of 10 are skeptical that the actions they take to protect their privacy will make any difference [41]. These issues of mistrust can persist because existing GoTCHAs take what Corbett and Le Dantec [18] argue is a "traditional" HCI approach of defining trust as a cognitive-based decision based on transparency and information—e.g., offering detailed timelines and accessible open data to instill trust—when a more nuanced, relational approach to trust is necessary with civic technology [28].

### 3.2 Design Approach

To support a more nuanced, relational approach to trust in GoTCHAs, specifically with regards to privacy and OBA, we propose the following research questions:

**RQ1 What are design requirements for people to participate in collective civic reporting of privacy harms from OBA?** *In other words, what concrete user interface or user experience characteristics are necessary to support this system?*

**RQ2 What cultural ideals do people have around redress for reported privacy harms from OBA, and the parties that carry out this redress?** *For example, how might the values*

*and beliefs that people hold about society shape their expectations of a relationship with the government and tech companies? How can these expectations surface additional considerations for the design of the system?*

Because the senses of resignation, powerlessness, and pessimism that people can have about GoTCHas and OBA can feel nearly intractable, we employ a mixture of design fiction methods to answer the above RQs, in hopes of encouraging people to feel more creative about the possibilities of GoTCHa design. As Bleecker [8] first introduced, fiction can “be a purposeful, deliberate, direct participant in the practices of science fact” that allows us to understand, explore, and question alternate futures. Within the context of user security and privacy, especially in challenging the hegemony of large tech institutions, there has already been some design fiction work. For example, Wong et al. [60] employed design workbooks—collections of conceptual designs—to first explore questions about user privacy stemming from a science fiction novel. More recently, Møller et al. [33] also explored the lack of agency of unemployed individuals over consent to data sharing through a fictitious job placement app that collects highly invasive personal data.

In the realm of design fiction, we employed a combination of three participatory design approaches in our study: **fictional inquiry**, **story completion**, and **comicboarding**. As an immediate descendent of design fiction, *fictional inquiry* is the practice of using partially fictional settings, artifacts, and circumstances to create a space for conducting collaborative design activities [22]. In this space, people are urged to imagine desirable futures and consider their everyday impacts. *Story completion* draws parallels with design fiction, asking participants to write or complete a fictional story about given a hypothetical seed scenario [17]. With the story completion method (SCM), people can share their thoughts about an idea without the burden of imagining or inserting themselves into the situation, allowing them to be more imaginative with the possibilities of the design, and be encouraged to think outside of immediate impacts on their own lives to about how different scenarios affect other people. Within HCI, SCM has been employed to identify potential thematic futures around human-VR pornography [61], human-robot [14], and human-AI voice assistant [11] interactions. To support these two approaches, we adapted *comicboarding*, a co-design method that provides a structure of comic strips and partially-completed content as a scaffold for users to come up with novel ideas. Comicboarding can be especially helpful in cases where people are not accustomed to brainstorming [43], e.g., when users feel powerless about their privacy.

In this work, we created two comicboards illustrating a fictional government tool that supports collective civic reporting of privacy harms from OBA, with an empty panel in each board, where participants could write fictional stories about the tool. We hoped this format could sufficiently introduce complex technical concepts, but also be open-ended enough to elicit rich design requirements and cultural ideals for both this fictional tool and its impacts.

### 3.3 Comicboard Development

We began our comicboard development by brainstorming textual descriptions of the ways users could interact with such a tool, eventually consolidating these descriptions into two key roles that lay-users could play in the life-cycle of a collective civic harm-reporting system. We loosely based these roles on a composite of the first few stages of the data management life cycle [5], such as data generation and collection, as well as cleaning and processing. We chose to focus on these roles as they were more likely to be accessible to regular users or easily teachable, as opposed to roles like data analysis, interpretation, storage, and sharing, which might require additional levels of technical expertise and access.

We viewed the first role, data generation and collection, as a reflection of existing GoTCHas that we noted in the previous subsection. The *fictional* component of this role is that the complaints filed

are related to *privacy harms, specifically in the context of OBA*, rather than other existing supported violations in other domains. We propose a second role of data cleaning and processing as a natural augmentation of GoTCHas. Platforms that support crowd-worker-based annotation tasks—e.g., Amazon Mechanical Turk and Prolific—are already prevalent in the research community, but have not been widely explored as a way to increase citizen engagement with GoTCHas. Prior work [63] has suggested that giving people a small stake in a collective process of voicing demands—e.g., through responding to others’ concerns and quality-checking others’ contributions—in response to privacy harms can instill a sense of empathy in them for strangers, and make them feel more impassioned about the subject. The broad fictional design that we will explore, then, is the linkage of these two roles within one new system.

We initially described these roles through short text summaries, following Jin et al.’s [36] recommendations for low-cost privacy assessment methods; however, through initial pilot testing, we found that pilot participants were quickly bored or overwhelmed due to having to keep track of a number of stakeholders and technical concepts. Through a series of discussions as a research team, we pivoted to create rough drawings of how fictional characters might inhabit these roles, as well as short captions for these sketches. With additional rounds of pilot-testing, we further refined the content of these drawings and captions so that they would lead to responses relevant to our research questions, while remaining both easily legible, and open-ended enough for rich responses.

A key tension in our design process, as aforementioned, was balancing the dense technical nature of concepts like algorithmic ad delivery systems, privacy harms, and responsible regulatory bodies, with our goal of eliciting creative, imaginative stories from our participants. For example, we eliminated references to official government agencies like the FTC, because they required too much exposition and could confuse and overwhelm participants. Instead, we referenced “the government” broadly, to evoke an institutional authority that could lend the tool legitimacy.

### 3.4 Comicboard Content

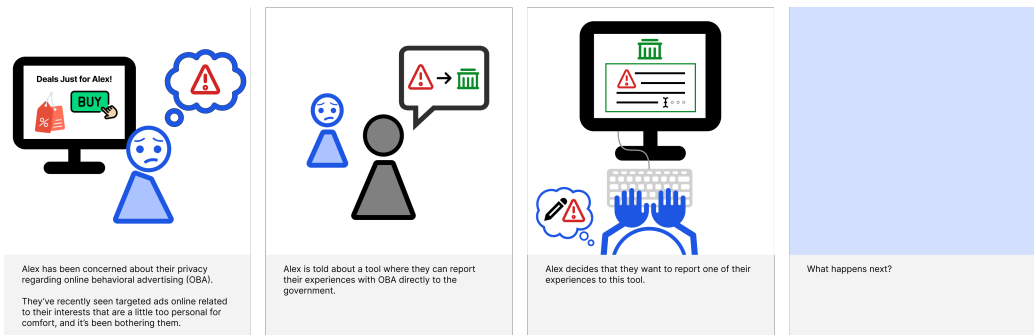
We developed two sets of comicboards, each with four panels, that describe the two potential roles that lay-users could play in the life cycle of a collective civic harm-reporting system.

The first set of panels (herein referred to as the “contribution” board) addresses the role of data generation and collection. Specifically, it follows the story of a character named Alex, who is concerned about their privacy in relation to OBA, and is told about a tool from the government where they can report their violating experiences with it. Subsequently, Alex decides they want to contribute a report of their experiences to this tool. The final fourth panel does not contain an image, and asks participants the question, “What happens next?” We left the *final* panel empty because we wanted participants to consider all the unknown possibilities for Alex’s contribution.

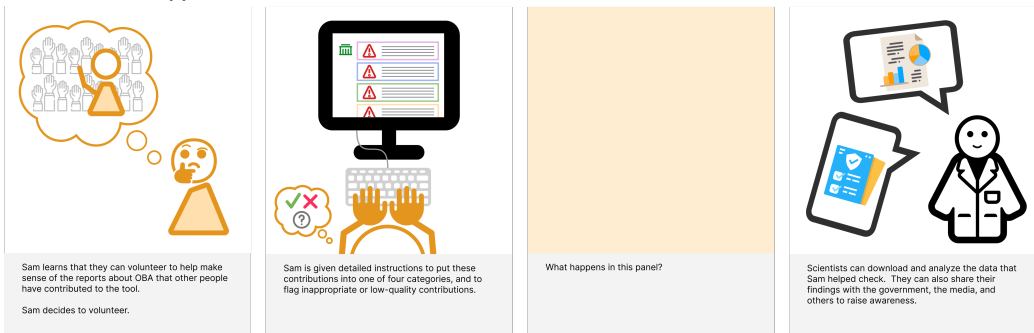
The second set of panels (herein referred to as the “annotation” board) covers the role of data cleaning and processing. Set in the same fictional world as the contribution board, the annotation board follows the story of Sam, who learns they can volunteer to make sense of the reports that other people have contributed to the tool. Sam decides to become a volunteer. The third panel does not contain an image, and asks participants the question, “What happens in this panel?” The final fourth panel explains that scientists can access the data the Sam helps annotate and clean, and can use it to create analyses and reports that will be shared with others to raise awareness. In this orange board, we purposefully left a panel empty in the *middle* of the story because we wanted participants to consider how Sam’s volunteering could be directly tied to the work of privacy experts and policy professionals, and how they fit into a process with many stakeholders.

Both contribution and annotation boards can be viewed in Figure 1.





(a) Panel 1: Alex has been concerned about their privacy regarding online behavioral advertising (OBA). They’ve recently seen targeted ads online related to their interests that are a little too personal for comfort, and it’s been bothering them.  
 Panel 2: Alex is told about a tool where they can report their experiences with OBA directly to the government.  
 Panel 3: Alex decides that they want to report one of their experiences to this tool.  
 Panel 4: What happens next?



(b) Panel 1: Sam learns that they can volunteer to help make sense of the reports about OBA that other people have contributed to the tool. Sam decides to volunteer.  
 Panel 2: Same is given detailed instructions to put these contributions into one of four categories, and to flag inappropriate or low-quality contributions.  
 Panel 3: What happens in this panel?  
 Panel 4: Scientists can download and analyze the data that Sam helped check. They can also share their findings with the government, the media, and others to raise awareness.

**Fig. 1.** The final two comicboards we developed. We refer to the first comicboard throughout the paper as the “contribution” board, and the second as the “annotation” board.

## 4 Study Design

We presented our comicboards in an online survey instrument deployed on Qualtrics and shared on Prolific, a crowd-work platform. We recruited a total of 50 participants for the main study to share design requirements, cultural ideals, and fictional stories about collective civic reporting of OBA harms.

### 4.1 Recruitment, Ethics, and Compensation

We first screened 200 people for eligibility for the main study, to ensure what we had a large enough pool of potential participants. We asked them about prior awareness of or familiarity with OBA. This was to avoid overloading participants with more new technical information about OBA as they

were being tasked with writing fictional stories, since the comicboards themselves were already dense with new ideas. These participants were also asked to confirm the demographic information they had previously provided to Prolific when joining the platform—that they were adults located in the United States and fluent in English. (There were no other exclusion criteria). This screener took on average less than a minute to complete, and participants were compensated 0.25 USD on the Prolific platform.

Out of the participants who indicated awareness of OBA, 50 participated in our main study on a first-come-first-serve basis until saturation, reflecting sample sizes in similar prior work [11, 17, 61]. (Participant demographics are reported in Table 1). The main survey took on average 19 minutes and 33 seconds to complete, for which participants were compensated 9 USD. Our study was approved by an institutional review board.

## 4.2 Survey Instrument

There were four main components to the survey instrument: a comprehension check, the first contribution comicboard, the second annotation comicboard, and a summary of all comicboards.

First, we provided participants with detailed instructions on the context of the study, explaining that they were about to read and see images related to a fictional world, and that they would be asked to write stories about this fictional world. We also explained that participants could be as creative as they wanted, and there were no right or wrong ways to write these stories. We then asked participants to confirm their understanding of these instructions.

Then, we showed participants the contribution comicboard (Figure 1), involving a fictional character, Alex, and asked them to share any features or information they felt was necessary for Alex to make a successful contribution to the government tool. We then asked participants about any hopes, concerns, or other thoughts Alex had about the contribution process or the tool itself. Finally, we asked participants to write a short story completing the empty panel in the comicboard. We placed the story completion task after the first two questions to prime participants to include richer details in their final stories.

Third, we showed participants the annotation comicboard, involving a second fictional character, Sam (again, see Figure 1). Participants were asked about any obstacles Sam might encounter in annotating others' contributions, as well as what Sam thought about the other people who've interacted with the tool or will do so in the future. Then, once again, we asked participants to write a short story completing the empty panel in the comicboard.

Finally, we showed participants all of the comicboards they had seen previously in the study, as well as all of the responses and stories they had written. Participants were asked about the extent to which they would like to live in the fictional world they wrote about, and to explain why or why not, similar to past work [11, 61].

## 4.3 Analysis

Systems supporting user privacy, especially in the context of privacy harms, always have multiple, interconnected stakeholders and involve countless intertwined decisions. Our comicboards themselves were self-referential and inherently related, and we also asked participants to reflect across all comicboards. Consequently, similar to prior comicboarding and storyboarding studies in HCI [32, 38], we intentionally analyzed participants' responses *across* storyboards and questions, rather than analyzing specific responses to each question or per story written. We did not actively screen for use of large language models by participants, as there are no foolproof or reliable methods for doing so; however, we did not come across any responses that were low-effort or low-quality in our analysis.

| Demographic       | Group                           | n         |
|-------------------|---------------------------------|-----------|
| Age               | 18-24                           | 10        |
|                   | 25-34                           | 21        |
|                   | 35-44                           | 14        |
|                   | 45-54                           | 2         |
|                   | 55-64                           | 2         |
| Gender Identity   | Female                          | 24        |
|                   | Male                            | 26        |
| Ethnicity         | White                           | 41        |
|                   | Black                           | 5         |
|                   | Asian                           | 3         |
|                   | Mixed                           | 1         |
| Employment Status | Employed Full-Time              | 17        |
|                   | Employed Part-Time              | 4         |
|                   | Student                         | 8         |
|                   | Unemployed                      | 4         |
|                   | Homemaker, retired, or disabled | 3         |
|                   | Not reported                    | 19        |
| <b>Total</b>      |                                 | <b>50</b> |

**Table 1.** Demographics of Prolific participants in the main study.

We adopted a reflexive approach to thematic analysis, in that we consistently iterated upon and modified codes to purposefully construct higher-level themes—or as Braun and Clarke put it, a “continual bending back on oneself...questioning and querying the assumptions we are making in interpreting and coding the data” [9]. Two members of the research team first conducted open coding on a sample of 20 participants’ responses and generated a total of 75 initial codes. Simultaneously, we had started to develop 10 potential higher-order themes. The two researchers discussed disagreements and attempted to consolidate codes, ultimately agreeing that at this stage, we needed more granularity, and settling on a codebook of 19 codes. We applied this codebook to the full dataset of responses. Finally, we generated five final themes from these codes during repeated discussions with the rest of the research team. We present these five themes in the following section.

## 5 Study Findings

We group contents of participants’ responses into five key themes. The first theme, the **UI or UX elements** of the tool, refers to stories that referenced the visual elements of the tool, specific types of evidence that should be supported or ingested by the tool, information about the tool or the submission generally, and any other features participants felt the tool should have. The second theme, **post-contribution expectations**, deals with stories that mentioned Alex’s ideals for their experience after contributing a report, including speed, transparency, and accountability, as well as specific design elements related to those factors. The third theme, **costs of contributing**, covers stories that mentioned downsides of the tool entailed by both contributing and annotating; the fourth theme, **benefits of contributing**, deals with the opposite. Finally, the fifth theme, **outlook on the future**, explores participants’ optimism or pessimism about the tool and its place in society.

## 5.1 UI and UX Elements of Tool

A vast majority (N=40) of participants mentioned specific evidence of harm that should be supported or ingested by the tool in their stories and responses, which included both details about the specific harmful targeted ads, as well as *why* those ads were harmful. For example, as a start, as P43 summarized, “Alex would need to be able to provide images of the advertising he is getting as well as images or just information about his...personal life and how this advertising is infringing on it.” Participants also brought up direct links to the ads, as well as information about both the advertiser and the platform where the ad was served. Finally, some participants also noted that Alex would enter personal information about themselves, such as a Social Security Number (e.g., P36), browsing and search history (e.g., P23), and contact information so they could receive updates about how their contribution was being processed (more on this in the following subsection).

To support all these types of information, some participants (N=10) mentioned specific visual elements of the tool, including the method for delivering their contribution and associated evidence. For example, while P47 mentions that Alex would submit their contribution by email, other participants suggested that Alex would type into an online form with a large longform text box. Participants also mentioned other modalities for submissions, such as a mix of both multiple choice and free responses to ensure uniform data, as well as places to upload images, along with “*basic features found in any word processing application*” (P7). For the actual process of writing up a contribution, participants hoped Alex would be educated about the proper terminology to use in their contribution (P21), guidelines on word length and level of detail (P33), as well as tips for what Alex should do in the meantime while their contribution was being processed (P44).

However, a few participants also noted that the tool, as built by the government, would likely have poor documentation and guidance for appropriate submissions: “*There wasn’t a lot of detail on the website to guide [Alex] through the process*” (P14). As a similar dig toward government-based technology, P21 wrote,

*“Alex went to the website and was surprised to find that the front page of it looked fairly modern. However, when they went any deeper than the front page, the website looked like it was designed in 2002. They were able to figure out how to lodge a complaint, but they are not confident that it will accomplish the goal set forth. There was a lot of legalese on the website. It made it confusing at times, and they are not sure that they got their point across entirely.”* –P21

Finally, a few participants also expected that the tool would provide additional features, such as periodic security and privacy tips (P10, P18, P35), background monitoring of the ads Alex saw, ad-blocking capabilities integrated directly into the tool (P20, P37), as well as information about local community groups Alex can join for more involvement (P25).

## 5.2 Post-Contribution Expectations

Several participants (N=33) directly wrote in their stories about Alex’s expectations for the tool and the government after they had made their contributions, primarily focused on three qualities: speed, transparency, and accountability.

Eleven participants referenced some sort of **speed** element (or lack thereof) related to Alex’s expectations of feedback from the tool, ranging from immediate feedback; to waiting a few weeks, months, or years; to never hearing back at all. Those that wrote about Alex experiencing longer wait times or never hearing back also mentioned that while Alex expected this from the government, they were still disappointed and resigned. For example, as P41 wrote, “*Nothing seemed to come of his report. Two months later Alex noticed really no difference and just gave up on the whole thing,*

realizing it doesn't matter and everybody already has all his information anyways." And, as P19 succinctly wrote, "Alex couldn't believe he waited 6 months for nothing."

Relatedly, participants also expressed concern about a lack of **transparency** about what would happen to their contribution after they finished submitting it. Specifically, 8 participants brought up concerns over these contributions being collected for nefarious purposes, coupled with a lack of reassurance from the government. For example, P5 noted that Alex might not feel confident submitting so much personal information to a government website, and had "feelings of insecurity over whether or not the tool itself is going to invade their privacy in some way"; similarly, P45 wrote that Alex might fear retaliation for reporting, and did not have enough information to feel assured.

In a similar vein, participants wrote about a lack of **accountability** from the government after Alex submitted their contribution, particularly via their perceptions of the government as impersonal. For one, a few participants felt concerned that Alex's contribution would be dismissed since the harms Alex's reported from OBA were not severe enough to warrant further investigation. As a step further, both P16 and P4 mentioned that the government might use opaque automated tools to filter out contributions, so Alex's contribution might not even be read by a human. P16 also noted that there was little recourse for this automated decision, and Alex couldn't appeal the government's decision to dismiss their contribution.

Participants' expectations of a lack of speed, transparency, and accountability from the tool and the government also presented as an overall skepticism of the government's abilities to demonstrate concrete help. As P44 wrote, "The government is very large and is known for moving slowly; though they might be working as fast as they can, it might not feel that way to Alex." As an escalation, participants also wrote about how a negative experience could influence Alex to permanently distrust the government: "[The tool suggests] to run some safety programs to help protect [Alex's] data, but no word is given on taking action against the advertising companies. Frustrated, Alex follows the suggestion but vows to never trust the government to help his problems again" (P14).

As a solution to these concerns, several participants highlighted the importance of "a place on the tool to leave their information so they can be gotten back to by whoever runs the tool" (P41), i.e., personal contact information for further follow up. Participants mentioned that the tool should communicate to Alex an expected timeline for receiving a response from the government, and include features such as ticket numbers and progress bars to provide transparency. To support accountability, a few participants also mentioned that being provided with the contact information of an actual human government representative would make them feel more at ease about the legitimacy of the tool.

### 5.3 Costs of Contributing and Volunteering

A large majority of participants (N=43) also mentioned costs or downsides associated with contributing and volunteering. These costs were largely associated with the annotation role, possibly because our participants were crowd-workers themselves and were especially attuned to the struggles that online research participants might face. From the perspective of Sam, the volunteer annotator character in the orange board, participants (N=42) primarily wrote about two costs: volunteer fatigue and uncertainty over correct annotation.

Fatigue from crowd work and the psychological harms of online content moderation have been well-documented in existing literature [2, 30], as has the power imbalance between the researchers who solicit annotations and the annotators themselves [39]. Our participants gave responses that mirror these prior findings; P1 paints a particularly grim picture of Sam's life:

*"Sam doesn't know why he volunteered for this, this is more difficult than he thought—he wished this was a paid position. The scientists keep the volunteers locked in a office until*

*they complete so many reviews. Sam is trying to find a way out so he can escape this weird and terrible place. Sam vows to never volunteer for anything ever again, he will only be paid for his work.”*—P1

Participants also frequently brought up the banality of the task, and how low quality or incomprehensible contributions could make it difficult for Sam to concentrate on the task. P25 also outlines two challenges that Sam might face as annotator:

*“The amount of time that it takes: it ends up becoming a major drain on him due to the amount of time that he invests into it. The toll that it takes flagging inappropriate content: the things that he sees are scarring and end up causing problems for him.”*—P25

The nuanced nature of these contributions also meant that some participants (N=21) felt Sam would feel pressured to annotate and check everything perfectly, and would be worried about making a mistake, especially if Sam is not given adequate training and instructions. For example, P28 writes: *“picking out the good contributions is no problem, but the ones on the fence of helpful or not could be a tough choice....[Sam] could go home concerned with some of his choices.”*

The potentially sensitive content of these contributions also drew mentions of costs from the perspectives of both Alex and Sam. As mentioned in the previous subsection, participants felt that Alex might be concerned about contributing private information that might be misused for nefarious purposes (especially if exposed in a data breach and the consequent risks, as P47 notes). And, on behalf of Sam, participants noted there might be pressure to maintain the privacy of the contributions they read:

*“Sam also needs to balance the need for user privacy with the requirements of a proper investigation, which can be a delicate and complex task. Lastly, addressing concerns and offering follow-ups to submitters who wish to remain anonymous presents its own set of challenges, as contact options may be limited.”*—P45

Finally, on a more selfish level, some participants felt that Sam simply might not want to know that much information about other people’s lives. For example, P33 writes that there is so much detailed personal information in the contributions that Sam encounters that they feel like they are invading others’ privacy just by reading, let alone annotating.

#### 5.4 Benefits of Contributing and Volunteering

While most participants (N=42) mentioned some sort of perceived benefit for either Alex or Sam in contributing and volunteering, these benefits were largely abstract. Participants noted both collective abstract benefits (N=30) and individual ones (N=21) for Alex and Sam. Collective benefits that participants mentioned included a sense of community with others, broad data privacy laws that support user control, and a general wish for “collective good”, rooted in an ideal that the government will listen to the people and have their interests in mind. In particular, being able to read others’ contributions, as Sam does, injects a sense of empathy in Sam:

*“[Sam] comes across Alex’s submission and really can tell that Alex is disturbed by how personal his ad targeting is. Sam sorts the data as instructed, hoping that Alex (and everyone else who submitted) will get some kind of closure with the issue, because the ad targeting has just gotten out of hand.”*—P49

Individual benefits comprised of a general sense of “doing the right thing” and feeling good about oneself for helping others. A few participants also envisioned a future where Alex becomes a privacy rights activist, seeded by Alex’s initial contribution to the tool. For example, as P1 writes, *“He doesn’t know it yet, but reporting this one website sets him on a path to get involved in politics and laws in the future. He also gets involved with disputes of AI in the online world.”*

Other participants wrote about the sense of satisfaction that Sam might derive from putting in hours of work volunteering: *“After a few more hours of work, the pride she felt when she sorted her last submission was incomparable. She knew she did something that mattered, and that made her happy”* (P42). While this satisfaction and pride was largely associated with simply volunteering at all—e.g., *“Sam feels he has a duty to help the people”* (P49)—a few participants characterized Sam as finding new purpose in their role as a volunteer, and being extremely dedicated to optimizing and improving the annotation process. For example, P14 described Sam’s elation after devoting hundreds of hours to building an automated annotation tool: *“These long hours seemed to have finally paid off, with the most recent pass-through generating a nearly 99% success rate!”*

Only 15 participants mentioned possible concrete or immediately tangible benefits from the tool. These primarily consisted of seeing fewer or zero ads or monetary compensation for exposure to harmful ads. While a few participants constructed a fictional world where Sam was compensated for annotation, either through money (P27) or small trinkets (P36), several participants included some variation of *“Sam wishes they were being paid for annotation”* in their stories, evoking aforementioned costs associated with volunteer fatigue.

## 5.5 Outlook on Future of Tool and OBA

Beyond desired or anticipated costs and benefits of the tool, participants made several meta-level observations about the future of the tool, including Alex and Sam’s views on other contributors, and general outlooks on society.

Fewer than half of participants (N=22) felt optimistic about the future of the tool and its place in society; only 18 participants responded that they would like to live in the fictional world they had written about. These participants shared that they would appreciate a world in which the government listens to their privacy concerns and takes actions in response. Echoing the previous subsection, participants also expressed gratitude toward other people who took the time and effort to care for each other, and noted that this gratitude could be a motivation for continuous contribution to the tool. As P32 remarks, *“the people in the stories are much more considerate than people are in real life.”* P13 hypothesized that in the future, there would be fewer contributors to the tool because there would be fewer problems to report, due to the success of the tool.

Overall, however, participants tended to express pessimism about the tool’s abilities to change anything in response to OBA harms, especially from Sam’s perspective after reading through others’ contributions. In particular, many participants felt that reading others’ contributions made Sam feel even *worse* about the future of OBA harms, since there was no evidence of positive changes. For example, as P16 exemplifies:

*“Sam begins with a lot of gusto and determination but soon finds that the complaints are all quite similar. Some are superficial and of not much concern but others are extremely personal and Sam feels embarrassed for these people he is learning about. He really wants to help them but he realizes that the same type of complaints that he submitted are submitted by the thousands and nothing seems to be changing. He realizes that this is just another way for those in power to get more information from everyday people to use against them.”*—P16

Relatedly, participants also expressed negative judgments about people who had contributed to the tool. For example, some participants felt that people who contributed reports were delusional and wasting their time: *“Sam thinks people are generally using the tool like its going to cure a virus. That’s not the case at all and Sam feels bad for those people”* (P41). These negative perception extended to participants’ views of the annotators as well, who might purposefully submit low-quality annotations because they felt contributors were unintelligent. For example, P27, who wrote

| Theme                          | RQ1: Design Requirements   | RQ2: Cultural Ideals   |
|--------------------------------|--|--|
| UI and UX elements of tool     | <ul style="list-style-type: none"> <li>• Ingesting details related to specific harmful ads, rather than OBA as a generally harmful phenomenon</li> <li>• Detailed submission guidelines</li> </ul>                   | <ul style="list-style-type: none"> <li>• Interim assistance while contributions are being processed (e.g., security and privacy tips)</li> </ul>   |
| Post-contribution expectations | <ul style="list-style-type: none"> <li>• Evidence of progress or success</li> <li>• Form for leaving contact information</li> <li>• Timelines, ticket numbers, progress bars</li> </ul>                              | <ul style="list-style-type: none"> <li>• Speed</li> <li>• Transparency</li> <li>• Accountability</li> <li>• A human reading and responding to contributions</li> </ul>                                     |
| Costs of contributing          | <ul style="list-style-type: none"> <li>• Protections for volunteers (e.g., fatigue, sensitive contributions)</li> <li>• Protections for interpersonal privacy risks (e.g., reading others' contributions)</li> </ul> | <ul style="list-style-type: none"> <li>• Trade-offs between requiring highly personal information and convincing people to contribute</li> <li>• Grassroots versus government oversight of tool</li> </ul> |
| Benefits of contributing       | <ul style="list-style-type: none"> <li>• Immediate compensation (e.g., money) or tangible benefits (e.g., a coffee mug)</li> </ul>   | <ul style="list-style-type: none"> <li>• Duty to help others</li> <li>• Pride in volunteering</li> <li>• Sense of community</li> </ul>   |
| Outlook on the future          | <ul style="list-style-type: none"> <li>• Insurance against bad actors within collective of contributors</li> </ul>   | <ul style="list-style-type: none"> <li>• Privacy harms from OBA are only secondary concerns</li> </ul>   |

**Table 2.** Summary of qualitative findings as related to our research questions.

that Sam thought contributors were “*idiots for complaining*”, also described Sam as uncaring and motivated solely by monetary compensation to complete their annotations:

*“Sam decides to randomly categorize the results, using an ad hoc algorithm based on the first letter of the person’s complaint. Since he is paid per complaint categorized, he is able to process the complaints quickly, and make much more money than if he completed the task as intended.”*—P27

These negative outlooks were compounded by a handful of participants (N=7) who felt that the fictional world represented in the comicboards was boring and incomplete, as “*the only thing this world seems to focus on is internet privacy and OBA. I do believe these topics are important, but I want to live in a world where there is more than that*” (P46). As P6 wrote bleakly, “*There’s nothing but computers and scientists.*”

## 5.6 Overall Findings

We summarize our overall findings in references to our research questions below. Table 2 also presents how our RQs relate to the five qualitative themes we detailed above.



**5.6.1 RQ1. What are design requirements for people to participate in collective civic reporting of privacy harms from OBA?** Participants had clear conceptions of what a fictional tool for privacy harms from OBA would look like, heavily guided by their prior experiences interacting with government websites. They envisioned support for multiple kinds of ad-specific data, detailed guidelines for submission, and clear communication about feedback and claim statuses; along the way, they wanted assurance for protecting for sensitive personal information. Participants also highlighted a need for greater protections and benefits for volunteers (or perhaps doing away with the volunteer nature of the tool altogether): they frequently mentioned immediate compensation or tangible benefits directly tied to contributing or annotating, as well as support for volunteer fatigue. Finally, participants wanted internal safeguards against bad actors who might disrupt or slow down the reporting process.

**5.6.2 RQ2. What cultural ideals do people have around redress for reported privacy harms from OBA, and the parties that carry out this redress?** While participants noted that privacy harms from OBA were not necessarily at the forefront of their concerns, they also felt that the tool could inspire a sense of community and duty to help others around them, especially via the opportunity to read through others' reports. Consequently, people might find pride in volunteering their time and efforts to the tool and feel encouraged to participate in privacy activism more broadly. However, these abstract positive feelings alone might not fully offset potential costs of contributing: in particular, participants worried that the fictional tool could just be another way for their personal data and labor to be exploited by the government. Relatedly, participants expressed clear expectations for how they hoped the government would treat their concerns: they wanted speedy but personal responses that did not leave them feeling like just another number in a pile of reports.

## 6 Design Principles for Instilling Trust in GoTCHas

While our design fictions were specifically tied to online behavioral advertising and privacy harms, participants frequently spoke about their ideals for the government's operational capacity to help people more broadly when referring to the fictional tool. Their responses give us valuable insights into what people think about reporting and processing not just *privacy* harms, but also *consumer* harms more broadly.

In this section, we present seven design principles for instilling trust in GoTCHAs, derived from the five themes that we found in our analysis: **(1) visible, upfront benefits; (2) timely, useful feedback; (3) contestability; (4) error prevention measures; (5) integration into everyday life; (6) consideration of social influence; and (7) commitment diversity and flexibility.** These principles are scoped for the design of not only UI and UX components, but also government and societal processes and relationships. For example, principles 2 and 4 draw upon a combination of several of Nielsen's 10 usability heuristics for user interface design [44]. Nielsen's heuristics, widely regarded in the human-computer interaction field as a canonical reference for low-cost evaluation, reflected several of our RQ1 findings, which might arise when people directly and actively interact with a GoTCHa. The remaining principles, more rooted in RQ2, can be used to describe conditions *outside* of the tool that need to be met so that GoTCHAs can be fully successful and trustworthy—hence why existing GoTCHAs tend to only partially satisfy them.

In the following subsections, we define each of our principles. We then apply these principles in a preliminary evaluation of a limited sample of existing GoTCHAs—Regulations.gov, the FTC fraud reporting site, the FTC identity theft reporting site, and the respective complaint sites for the CFPB and FCC—summarized in Table 3. Examples of the principles being satisfied are shown in Table 4.

| Design Principle                              | Regulations.gov | FTC (Fraud) | FTC (Identity Theft) | CFPB | FCC |
|---|-----------------|-------------|----------------------|------|-----|
| <i>Visible, upfront benefits</i>              | ✗               | ✓†          | ✓†                   | ✓†   | ✓†  |
| <i>Timely, useful feedback*</i>               | ✗               | ✓†          | ✓†                   | ✓†   | ✗   |
| <i>Contestability</i>                         | ✗               | ✓†          | ✓†                   | ✓†   | ✓†  |
| <i>Error prevention measures*</i>             | ✗               | ✓           | ✓                    | ✓    | ✗   |
| <i>Integration into everyday life</i>         | ✗               | ✗           | ✓†                   | ✗    | ✗   |
| <i>Consideration of social influence</i>      | ✗               | ✗           | ✗                    | ✗    | ✗   |
| <i>Commitment diversity &amp; flexibility</i> | ✗               | ✗           | ✗                    | ✗    | ✗   |

**Table 3.** Initial evaluation of potential existing GoTCHAs using our design principles.

\*: Design principle is based on a Nielsen heuristic.

✓: Satisfies this design principle.

✗: Does not address this principle at all.


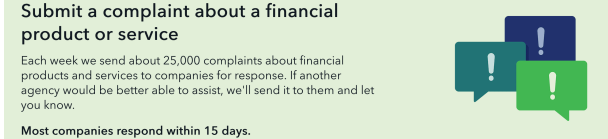
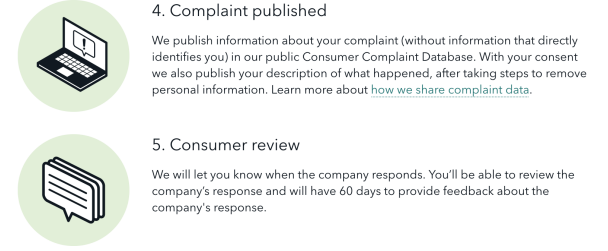
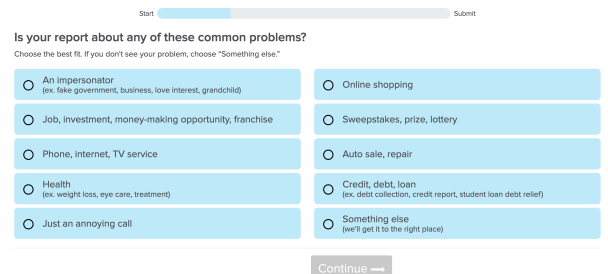
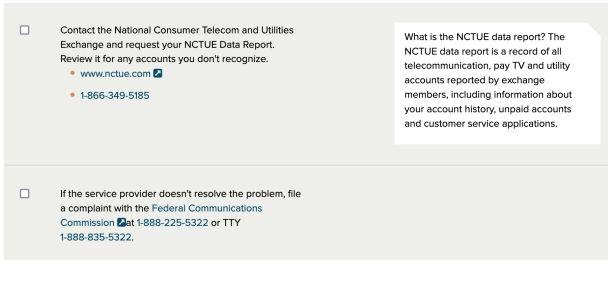
†: Partially satisfies this principle.

## 6.1 Visible, Upfront Benefits

Several participants noted that they would be unmotivated to report harms in real life, especially since it was difficult to see how their contributions would directly lead to a positive outcome or concrete benefit. The GoTCHAs we evaluated do little to strengthen this tenuous link. The FTC’s identity theft site, for example, clearly tells people the immediate next steps of reporting cases of identity theft: after submission, the system can create a case for the reporter, and the reporter will be given a checklist of actions they can take on their own while their case is being processed. However, it’s not clear what outcomes a reporter might expect from the process; while the checklist is helpful, it does not directly address any damages or policy changes people might expect from reporting. The FTC’s fraud site, on the other hand, informs people from the outset that they will not resolve their individual case, but hints that people might be able to recoup financial losses from falling prey to scams later on. These relationships—between reporting action and expected outcome—are further complicated by the aforementioned possibility of cases being delegated to third-parties who may or may not fully address people’s concerns. To motivate contributions, then, GoTCHAs should present and communicate **visible, upfront benefits** to people: examples might include descriptions of financial remuneration in prior resolved cases, discussions of how these contributions are being concretely used to put together policy briefs, or access to online tools and resources while contributors wait for a response.

## 6.2 Timely, Useful Feedback

Our participants highlighted the importance of hearing back about their contributions in a reasonable time frame as a bare-minimum requirement for their trust in the fictional tool’s ability to help people. A guaranteed quick response might get Alex to consider contributing, but they certainly wouldn’t get their hopes up for one, much less expect the response to be actually helpful. But the response needs to be not only quick, but also personal and informative about what exactly was happening with Alex’s contribution—not from an automated no-reply address. We liken this requirement to Nielsen’s [44] heuristic of “visibility of system status”—which argues that when users can understand a system’s state, they feel in control and trust in the system’s ability to do

| Design Principle   | Example Screenshot   |
|--|--|
| <p><i>Visible, upfront benefits</i></p> <p>The FTC Identity Theft site <b>partially satisfies</b> this principle. It identifies the immediate next steps of reporting and offers a checklist of actions to take while their case is being processed. However, the checklist does not make clear what outcomes (e.g., remuneration), if any, people might expect from the process.</p>  |  <p>IdentityTheft.gov can help you report and recover from identity theft.</p> <p><b>HERE'S HOW IT WORKS:</b></p> <p><b>Tell us what happened.</b><br/>We'll ask some questions about your situation. Tell us as much as you can.</p> <p><b>Get a recovery plan.</b><br/>We'll use that info to create a personal recovery plan.</p> <p><b>Put your plan into action.</b><br/>If you create an account, we'll walk you through each recovery step, update your plan as needed, track your progress, and pre-fill forms and letters for you.</p>  |
| <p><i>Timely, useful feedback</i></p> <p>The CFPB Complaints site <b>partially satisfies</b> this principle: it gives a rough timeline, but the response might be delegated to a third party not beholden to that time commitment.</p>   |  <p><b>Submit a complaint about a financial product or service</b></p> <p>Each week we send about 25,000 complaints about financial products and services to companies for response. If another agency would be better able to assist, we'll send it to them and let you know.</p> <p><b>Most companies respond within 15 days.</b></p>  |
| <p><i>Contestability</i></p> <p>The CFPB Complaints site <b>partially satisfies</b> this principle: it provides people one chance to respond at the end of the complaint process; however, there are no other chances for back-and-forth. Additionally, all GoTCHas sampled have links to privacy policies, but these policies have no room for interactability beyond consent.</p>  |  <p><b>4. Complaint published</b></p> <p>We publish information about your complaint (without information that directly identifies you) in our public Consumer Complaint Database. With your consent we also publish your description of what happened, after taking steps to remove personal information. Learn more about <a href="#">how we share complaint data</a>.</p> <p><b>5. Consumer review</b></p> <p>We will let you know when the company responds. You'll be able to review the company's response and will have 60 days to provide feedback about the company's response.</p>   |
| <p><i>Error prevention measures</i></p> <p>The FTC Scams website <b>fully satisfies</b> this principle by including step-by-step, clearly-defined questionnaires to standardize the types of complaints people can submit.</p>   |  <p>Start <span style="float: right;">Submit</span></p> <p><b>Is your report about any of these common problems?</b></p> <p>Choose the best fit. If you don't see your problem, choose "Something else".</p> <p><input type="radio"/> An impersonator (ex. fake government, business, love interest, grandchild)</p> <p><input type="radio"/> Online shopping</p> <p><input type="radio"/> Job, investment, money-making opportunity, franchise</p> <p><input type="radio"/> Sweepstakes, prize, lottery</p> <p><input type="radio"/> Phone, internet, TV service</p> <p><input type="radio"/> Auto sale, repair</p> <p><input type="radio"/> Health (ex. weight loss, eye care, treatment)</p> <p><input type="radio"/> Credit, debt, loan (ex. debt collection, credit report, student loan debt relief)</p> <p><input type="radio"/> Just an annoying call</p> <p><input type="radio"/> Something else (we'll get it to the right place)</p> <p><input type="button" value="Continue"/></p> |
| <p><i>Integration into everyday life</i></p> <p>The FTC Identity Theft site <b>partially satisfies</b> this principle: it provides people with a list of recovery steps in the event of potential identity theft, including listing out specific contact information and advice for speaking to other agencies, companies, and third parties who might provide more direct solutions. In doing so, it highlights for people the scope and severity of harms that identity theft has on everyday life. However, as a standalone website, people must intentionally go out of their way to visit and report harms.</p> |  <p><input type="checkbox"/> Contact the National Consumer Telecom and Utilities Exchange and request your NCTUE Data Report. Review it for any accounts you don't recognize.</p> <ul style="list-style-type: none"> <li>• <a href="http://www.nctue.com">www.nctue.com</a></li> <li>• 1-866-349-5185</li> </ul> <p><input type="checkbox"/> If the service provider doesn't resolve the problem, file a complaint with the Federal Communications Commission ☎at 1-888-225-5322 or TTY 1-888-835-5322.</p> <p><b>What is the NCTUE data report? The NCTUE data report is a record of all telecommunication, pay TV and utility accounts reported by exchange members, including information about your account history, unpaid accounts and customer service applications.</b></p>  |
| <p><i>Consideration of social influence</i></p>  | <p>No GoTCHas sampled satisfied this principle.</p>  |
| <p><i>Commitment diversity &amp; flexibility</i></p>   | <p>No GoTCHas sampled satisfied this principle.</p>  |

**Table 4.** Screenshot examples of existing GoTCHAs fully or partially satisfying our seven design principles. These screenshots were taken in October 2024.

what they ask it—and propose **timely, useful feedback** as a necessary characteristic of GoTCHAs. Currently, none of the GoTCHAs we evaluated fully satisfy this principle. While a few make commitments to timeliness (e.g., within 60 days) and inform people of what they might *expect* in a response, there are no guarantees of this timeline, and the response itself is sometimes delegated to third-parties (e.g., law enforcement, other government agencies, external companies) that might not be beholden to these commitments.

### 6.3 Contestability

Multiple participants expressed concerns about Alex receiving impersonal decisions from the government about their complaint that they could not appeal, e.g., “*an automatically generated email...saying that according to the standards of the government, the advertisements [Alex] was being shown were perfectly legal and there is nothing more they can do.*” Responses like these could further contribute to people feeling powerless to effect change, and may discourage them from contributing reports in the future. Giving people a space to respond to these potential risks and decisions can help them feel more confident in their choice to make contributions in the first place. Leaving room for **contestability** in GoTCHAs—inspired by Rakova et al. [46]—sets up an ongoing dialogue between people and the government and can signal a commitment on the GoTCHa’s part to pursue systemic changes, rather than ad hoc ones.

Relatedly, participants also worried about Alex’s contributions being used against them—i.e., if the detailed personal information in their contributions was used to target them in other ways. Participants forecasted a sense of betrayal and powerlessness if this were to happen. While it is nigh impossible to guarantee the security and privacy of a system [31], GoTCHAs can instill more trust in people by clarifying how they can be held accountable as data stewards. While all of the GoTCHAs in the sample link to their respective agencies’ privacy policies as a way to assure people that their privacy will be protected and their contributions will not be misused, there is little discussion of recourse if this data is violated in, e.g., a data breach. There is also little interactability with the privacy policy on the public’s part. Thus, none of the GoTCHAs we evaluated fully satisfy this design principle.

### 6.4 Error Prevention Measures

When participants wrote about Sam, the character in the annotator role, they frequently expressed how Sam could be frustrated by incomprehensible, irrelevant, or low-quality contributions. One immediate alleviation for this problem is **error prevention measures**—drawn from Nielsen [44], again—that clearly define what kind of contributions a GoTCHa can ingest, with formatting built into the tool itself. Both FTC websites and the CFPB satisfy this design principle by including step-by-step, clearly-defined questionnaires to standardize the types of complaints people can submit. This results, in the case of the CFPB, in a tidy database of complaints that is easily searchable and accessible to the public; such a structured and comprehensible data output can also aid contestability. The FCC, on the other hand, fails in this principle: submission formats vary widely based on the type of complaint (i.e., phone, Internet, television, etc.), and the submission sites are more akin to open-ended email composition boxes rather than professionally-organized intake forms, with neither guidance on what details people should include in their submissions, nor limits on word count or file types to attach. Regulations.gov is equally open-ended, with only a small text box, file upload link, and contact information field; there are no affordances for what people should contribute, or to stop them from commenting on the wrong topic.

## 6.5 Integration into Everyday Life

Currently, all the GoTCHas we evaluated operate as standalone websites that require people exit their daily lives to visit if they want to report a harm. In other words, it's unlikely that that people will report harms unless they are individually aware of these platforms and especially motivated to do so, e.g., only in severe cases. Simultaneously, the FTC itself can primarily only respond to these "flashier" events, rather than setting up sustained efforts to the more-ubiquitous but smaller harms, due to being perpetually underresourced [34]. However, people's unrelenting exposure to these "small" harms can mean that they become habituated or numb to its effects over time, accepting them as a fact of life [45]. One way we might address this is by making *harm recognition and reporting* as accessible and mundane as the harms themselves have become, through **integration into everyday life**. While the FTC Identity Theft reporting platform is not accessible outside of its eponymous website, it makes a concerted effort to teach reporters about all of the different domains of life identity theft can affect, including housing, education, banking, and social services. It also provides people with direct links and contact information for the external agencies that might help with these domains, along with the exact language to use when dealing with different agencies about identity theft, e.g., through email and letter templates.

Future GotCHa design might consider alternative modalities for encouraging harm recognition and reporting. According to a Statista poll<sup>2</sup>, over a quarter of users in the United States use an ad-blocking tool when browsing, which typically takes the form of an unobtrusive browser extension that runs in the background. We can imagine GoTCHAs being combined with such relevant existing tools, as our participants suggested. Or, the GoTCHAs themselves might run in the background in a similar manner and deliver periodic short questionnaires to people, perhaps through an Experience Sampling Method approach, taking a continuous "pulse" of harm measurements.

## 6.6 Consideration of Social Influence

Social influence can strongly affect people's motivation to act, especially when uncertain [15]. If it is clear to people that many others report OBA harms, then people may feel empowered to do the same. In contrast, if it is clear that others largely ignore these harms, people may feel disincentivized to file reports themselves. Participants mentioned both positive and negative opinions of the other people who might use the fictional tool, which in turn affected their outlook for the tool and desire to use it. Many participants in our survey expressed solidarity and empathy with other users of our hypothetical reporting tool, contributing to a sense of community, while others expressed pessimism on considering the harms of OBA or the perceived futility of relying on our reporting tool to address them. These reactions all prompt a necessary **consideration of social influence** when designing GoTCHAs, especially if the goal is to encourage and sustain civic participation. Currently, however, none of the systems in our sample do much to address this design principle beyond simply *having* a publicly accessible database of consumer complaints that people can read. These systems could instead encourage community building outside of the GoTCHa, perhaps through offering contact points of local offices and advocacy organizations, connecting people with support networks, or supporting new and existing online communities dedicated to addressing these issues.

## 6.7 Commitment Diversity and Flexibility

Currently, while GoTCHAs allow self-motivated individuals to access and download contribution data for personal use, the sole primary public-facing role that people can take on when interacting with these systems is that of a reporter or complainant. In other words, currently, if people want to

<sup>2</sup><https://www.statista.com/topics/3201/ad-blocking>

interface with GoTCHas, they can really only do so through filling out the respective web form and waiting for a response. However, our participants described both Alex and Sam's roles in the fictional with varying levels of dedication and responsibility. Alex might contribute a report and forget about it, or they might feel emboldened to become a privacy legislation activist and recruit others to become contributors. Sam might annotate contributions just out of boredom, or they might devote hundreds of hours toward programming a more efficient annotation system to improve the performance of the fictional tool. Or, Alex and Sam might vacillate between any of these spectra of activity, or even decide to seek employment with the government and become policy professionals themselves. Where appropriate, then, GoTCHas should *formally* provide their users with the opportunity for **commitment diversity and flexibility**; we envision a plethora of roles that people could play in these systems, such as not only recruiters and developers (as aforementioned), but also analysts, modelers, and contesters of contributions [47], as well as broadcasters of eventual findings. This allowance for different types of public participation may also allow motivated contributors to help under-resourced government agencies process mass volumes of reports more quickly.

## 7 Limitations and Additional Considerations

In this section, we summarize some limitations of the work, including our design fiction method and its impact on participant outlook, and our participant sample and study context. We also propose some considerations for future research in the design of GoTCHas: some potential pitfalls of privacy GoTCHas, the generalizability of GoTCHas to other sectors, and collective empowerment.

### 7.1 Limitations

*7.1.1 Participant Sample and Study Context.* Our participants were recruited from Prolific, an online crowdwork platform. It is thus not out of the question that they may have been uniquely sensitive and attuned to the issues of volunteer fatigue and exposure to harmful content echoed in prior work [2, 55]. Future work could explore the co-design of GoTCHas with crowd workers and content moderators, while heeding recommendations from Irani and Silberman's experiences with Turkopticon [35]. Additionally, our scope was limited to online behavioral advertising and GoTCHas within the United States of America. Future work could explore how a GoTCHa for privacy might look for other kinds of privacy harms, such as data breaches and workplace surveillance, or how concepts such as everyday algorithmic auditing [52] might be institutionalized. Examining systems in other regulatory environments or cultural contexts could also prove fruitful in the future.

*7.1.2 Design Fiction and Participant Outlook.* Participants in our study were, on average, pessimistic about the future of online behavioral advertising, and unenthusiastic about the capacity of a government tool to address OBA's harms. Perhaps OBA is already so prevalent and bleak that no single solution can inspire hope for change, or perhaps the overall harms of OBA are simply not striking enough for people to care about in a vacuum (again, recalling [45]). While both of these hypotheses might prove true, the setup of our comicboards themselves might have also encouraged more negative responses. For example, our comicboards were deliberately open-ended and sparse, because we specifically wanted participants to be as creative as possible in their story writing; however, this might have led them to fill in the blank with their existing perceptions of user privacy and the government, which generally trend negative in the United States [3]. In future work, we could consider seeding participants with examples of potential positive outcomes from GoTCHas, e.g., reparations or specific policy changes, in order to encourage more idealistic stories to which we can aspire through design.

## 7.2 Additional Considerations and Reflections

*7.2.1 Pitfalls of Privacy GoTCHAs.* Setting up a GoTCHa, especially for privacy harms, would require the utmost care in delineating ownership, authorship, and access. For example, some participants felt wary of associating too much personal information with their reports of harm, since they were concerned that the government could use their contributions to re-target them with even more invasive ads. Is there any party that users would trust enough to act as stewards over their contributions to the GoTCHa? How might user acceptability for granting access vary based on different parties? What does insurance against bad actors *within* the population of contributors look like? Future work should explore the tradeoffs associated with such access control, while paying mind to our design principles (in particular, contestability and consideration of social influence).

*7.2.2 Application to Non-Government Harm-Reporting Tools.* While in this work, we exclusively explored the concept of a *government*-hosted tool for citizen harm-reporting, we argue many of these design principles can apply to tools created by other parties as well. For example, the website Top Class Actions<sup>3</sup> aggregates details of ongoing class action lawsuits, including those related to privacy harms, such as data breaches and misuse of consumer data, and provides links for people to file claims and join classes. Adapting such platforms to be more in-situ—for example, by providing people with a periodic digest of class actions they might be interested in joining—could be a promising area of future impact, particularly in the face of off-slow-moving government changes.

*7.2.3 Collective Empowerment.* Much of our focus has been on the experiences of *individuals* who might contribute to or volunteer for a privacy GoTCHa—e.g., Alex and Sam—and it would be no surprise if people do find power and agency in reporting their privacy harms and receiving direct recompense for doing so. On the other hand, a persistent finding from our participants' stories was that people can view themselves as contributing to a larger movement over which they have collective ownership, rather than being solely motivated through protecting individual privacy. We might thus see privacy GoTCHAs, along with other crowd-powered algorithmic auditing methods, as a direct counter for not only combating the slow violence of privacy harms from OBA, but also challenging long-standing conceptions of privacy as something individuals act on only for themselves [49, 64]. As McDonald and Forte [42] write, "Individualism assumes that people have equal voice in articulating their privacy and defending it when we know that not to be the case."

Echoing Seberger et al. [48], a privacy GoTCHa in its current conception can only afford people a conditional empowerment over their privacy, with an endless supply of government agencies and tech companies ready to extract insights from harm reports. As designers and researchers, we have a duty to scaffold and nurture it into something collective and persistent, such that people can, over time, construct their own deeper understandings of the slow violence of privacy harms and engage with it on their own terms (recalling both "slow observation" [21] and "counter-data action" [19, 20]). Similar to how smartwatches and fitness trackers enable users to be actively engaged in their understandings of their physical health, while granting them the ability to communicate with their doctors with high-level concrete data, we might imagine that taking regular pulses of privacy harm reports can empower people to speak more definitively about their experiences—not only with professional privacy advocates and policy experts, but also with each other.

## 8 Conclusion

In this work, we explored—through a blend of fictional inquiry, story completion, and comicboarding—fictional futures in which a government-supported tool could facilitate people reporting on the privacy harms they experience from online behavioral advertising. Through an online survey, we

<sup>3</sup><https://topclassactions.com/>

found that participants had detailed conceptions of the user experience of such a tool, but wanted safeguards to prevent them from being exploited further for their data by the government itself. Consequently, participants also expressed a broad and deep distrust in the government's capacity to appropriately bring about mitigations for these harms. We extrapolated these design findings to existing government complaint-reporting tools in other domains, finding that they, too, lacked key qualities to instill trust; such systems are ripe for future design exploration using the design principles we propose as a starting point.

## References

- [1] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2019. Discrimination through optimization: How Facebook's Ad delivery can lead to biased outcomes. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–30.
- [2] Andrew Arsht and Daniel Etcovitch. 2018. The human cost of online content moderation. *Harvard Journal of Law and Technology* 2 (2018).
- [3] Brooke Auxier, Lee Rainie, Monica Anderson, Andrew Perrin, Madhu Kumar, and Erica Turner. 2019. Americans and privacy: Concerned, confused and feeling lack of control over their personal information. *Pew Research Center* (2019).
- [4] Tae Hyun Baek and Mariko Morimoto. 2012. Stay away from me. *Journal of advertising* 41, 1 (2012), 59–76.
- [5] Alexander Ball. 2012. Review of data management lifecycle models. (2012).
- [6] Colin J Bennett and Jesse Gordon. 2021. Understanding the “micro” in political micro-targeting: An analysis of Facebook digital advertising in the 2019 Federal Canadian election. *Canadian Journal of Communication* 46, 3 (2021), 431–459.
- [7] Colin J Bennett and David Lyon. 2019. Data-driven elections: Implications and challenges for democratic societies. *Internet policy review* 8, 4 (2019).
- [8] Julian Bleecker. 2022. Design fiction: A short essay on design, science, fact, and fiction. *Machine Learning and the City: Applications in Architecture and Urban Design* (2022), 561–578.
- [9] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative research in sport, exercise and health* 11, 4 (2019), 589–597.
- [10] Ryan Calo. 2011. The boundaries of privacy harm. *Ind. LJ* 86 (2011), 1131.
- [11] Julia Cambre, Samantha Reig, Queenie Kravitz, and Chinmay Kulkarni. 2020. "All Rise for the AI Director" Eliciting Possible Futures of Voice Technology through Story Completion. In *Proceedings of the 2020 ACM designing interactive systems conference*. 2051–2064.
- [12] Pew Research Center. 2022. Americans' views of government: Decades of distrust, enduring support for its role. *Pew Research Center* (2022).
- [13] Stevie Chancellor, Michael L Birnbaum, Eric D Caine, Vincent MB Silenzio, and Munmun De Choudhury. 2019. A taxonomy of ethical tensions in inferring mental health states from social media. In *Proceedings of the conference on fairness, accountability, and transparency*. 79–88.
- [14] EunJeong Cheon and Norman Makoto Su. 2018. Futuristic autobiographies: Weaving participant narratives to elicit values around robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 388–397.
- [15] Robert B Cialdini. 2009. *Influence: Science and practice*. Vol. 4. Pearson education Boston, MA.
- [16] Danielle Keats Citron and Daniel J Solove. 2022. Privacy harms. *BUL Rev.* 102 (2022), 793.
- [17] Victoria Clarke, Virginia Braun, Hannah Frith, and Naomi Moller. 2019. Editorial introduction to the special issue: Using story completion methods in qualitative research. , 20 pages.
- [18] Eric Corbett and Christopher Le Dantec. 2021. Designing civic technology with trust. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [19] Morgan Currie, Britt S Paris, Irene Pasquetto, and Jennifer Pierre. 2016. The conundrum of police officer-involved homicides: Counter-data in Los Angeles County. *Big Data & Society* 3, 2 (2016), 2053951716663566.
- [20] Craig Dalton and Jim Thatcher. 2014. What does a critical data studies look like, and why do we care? Seven points for a critical approach to 'big data'. *Society and Space* 29 (2014).
- [21] Thom Davies. 2018. Toxic space and time: Slow violence, necropolitics, and petrochemical pollution. *Annals of the American Association of Geographers* 108, 6 (2018), 1537–1553.
- [22] Christian Dindler and Ole Sejer Iversen. 2007. Fictional inquiry—design collaboration in a shared narrative space. *CoDesign* 3, 4 (2007), 213–234.
- [23] Liza Gak, Seyi Olojo, and Niloufar Salehi. 2022. The distressing ads that persist: Uncovering the harms of targeted weight-loss ads among users with histories of disordered eating. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–23.



- [24] Maximilian Golla, Miranda Wei, Juliette Hainline, Lydia Filipe, Markus Dürmuth, Elissa Redmiles, and Blase Ur. 2018. “What was that site doing with my Facebook password?” Designing Password-Reuse Notifications. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. 1549–1566.
- [25] Ayelet Gordon-Tapiero, Alexandra Wood, and Katrina Ligett. 2022. The Case for Establishing a Collective Perspective to Address the Harms of Platform Personalization. In *Proceedings of the 2022 Symposium on Computer Science and Law*. 119–130.
- [26] Ashley Gorski. 2022. The Biden Administration’s SIGINT Executive Order, part II. <https://www.justsecurity.org/83927/the-biden-administrations-sigint-executive-order-part-ii/>
- [27] Julia Hanson, Miranda Wei, Sophie Veys, Matthew Kugler, Lior Strahilevitz, and Blase Ur. 2020. Taking Data Out of Context to Hyper-Personalize Ads: Crowdworkers’ Privacy Perceptions and Decisions to Disclose Private Information. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [28] Mike Harding, Bran Knowles, Nigel Davies, and Mark Rouncefield. 2015. HCI, civic engagement & trust. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 2833–2842.
- [29] Sarah Hartmann, Agnes Mainka, and Wolfgang G Stock. 2017. Citizen relationship management in local governments: The potential of 311 for public service delivery. *Beyond bureaucracy: Towards sustainable governance informatisation* (2017), 337–353.
- [30] Kenji Hata, Ranjay Krishna, Li Fei-Fei, and Michael S Bernstein. 2017. A glimpse far into the future: Understanding long-term crowd worker quality. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 889–901.
- [31] Cormac Herley. 2016. Unfalsifiability of security claims. *Proceedings of the National Academy of Sciences* 113, 23 (2016), 6415–6420.
- [32] Alexis Hiniker, Kiley Sobel, and Bongshin Lee. 2017. Co-designing with preschoolers using fictional inquiry and comicboarding. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 5767–5772.
- [33] Naja Holten Møller, Trine Rask Nielsen, and Christopher Le Dantec. 2021. Work of the Unemployed: An inquiry into individuals’ experience of data usage in public services and possibilities for their agency. In *Designing Interactive Systems Conference 2021*. 438–448.
- [34] Alysa Z. Hutnik and Laura Riposo VanDruff. 2021. “Not Outgunned, Just Outmanned” (For Now): Senate Hearing on Privacy Law Addresses Under-resourced FTC. <https://www.kelleydrye.com/viewpoints/blogs/ad-law-access/not-outgunned-just-outmanned-for-now-senate-hearing-on-privacy-law-addresses-under-resourced-ftc>
- [35] Lilly C Irani and M Six Silberman. 2016. Stories We Tell About Labor: Turkooption and the Trouble with “Design”. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 4573–4586.
- [36] Haojian Jin, Hong Shen, Mayank Jain, Swarun Kumar, and Jason I Hong. 2021. Lean privacy review: Collecting users’ privacy concerns of data practices at a low cost. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 5 (2021), 1–55.
- [37] Douwe Korff. 2022. The Inadequacy of the October 2022 New US Presidential Executive Order on Enhancing Safeguards for United States Signals Intelligence Activities. Available at SSRN 4495169 (2022).
- [38] Tzu-Sheng Kuo, Hong Shen, Jisoo Geum, Nev Jones, Jason I Hong, Haiyi Zhu, and Kenneth Holstein. 2023. Understanding Frontline Workers’ and Unhoused Individuals’ Perspectives on AI Used in Homeless Services. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [39] Clément Le Ludec, Maxime Cornet, and Antonio A Casilli. 2023. The problem with annotation. Human labour and outsourcing between France and Madagascar. *Big Data & Society* 10, 2 (2023), 20539517231188723.
- [40] Peter Mayer, Yixin Zou, Florian Schaub, and Adam J Aviv. 2021. “Now I’m a bit angry:” Individuals’ Awareness, Perception, and Responses to Data Breaches that Affected Them. In *30th {USENIX} Security Symposium ({USENIX} Security 21)*.
- [41] Colleen McClain, Michelle Faverio, Monica Anderson, and Eugenie Park. 2023. How Americans View Data Privacy. *Pew Research Center* (2023).
- [42] Nora McDonald and Andrea Forte. 2020. The politics of privacy theories: Moving from norms to vulnerabilities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [43] Neema Moraveji, Jason Li, Jiarong Ding, Patrick O’Kelley, and Suze Woolf. 2007. Comicboarding: using comics as proxies for participatory design with children. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 1371–1374.
- [44] Jakob Nielsen. 1994. Enhancing the explanatory power of usability heuristics. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 152–158.
- [45] Rob Nixon. 2011. *Slow Violence and the Environmentalism of the Poor*. Harvard University Press.
- [46] Bogdana Rakova, Renee Shelby, and Megan Ma. 2023. Terms-we-serve-with: Five dimensions for anticipating and repairing algorithmic harm. *Big Data & Society* 10, 2 (2023), 20539517231211553.

- [47] Andrew R Schrock. 2016. Civic hacking as data activism and advocacy: A history from publicity to open government data. *New media & society* 18, 4 (2016), 581–599.
- [48] John S Seberger, Hyesun Choung, and Prabu David. 2023. Problematizing “Empowerment” in HCAI. In *IFIP Conference on Human-Computer Interaction*. Springer, 270–279.
- [49] John S Seberger, Marissel Llavore, Nicholas Nye Wyant, Irina Shklovski, and Sameer Patil. 2021. Empowering resignation: There’s an app for that. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [50] John S Seberger, Irina Shklovski, Emily Swiatek, and Sameer Patil. 2022. Still creepy after all these years: the normalization of affective discomfort in app use. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [51] Aaron Shaw, Haoqi Zhang, Andrés Monroy-Hernández, Sean Munson, Benjamin Mako Hill, Elizabeth Gerber, Peter Kinnaird, and Patrick Minder. 2014. Computer supported collective action. *Interactions* 21, 2 (2014), 74–77.
- [52] Hong Shen, Alicia DeVos, Motahare Eslami, and Kenneth Holstein. 2021. Everyday algorithm auditing: Understanding the power of everyday users in surfacing harmful algorithmic behaviors. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–29.
- [53] Irina Shklovski, Scott D Mainwaring, Halla Hrunn Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and creepiness in app space: Perceptions of privacy and mobile app use. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 2347–2356.
- [54] Edith G Smit, Guda Van Noort, and Hilde AM Voorveld. 2014. Understanding online behavioural advertising: User knowledge, privacy concerns and online coping behaviour in Europe. *Computers in human behavior* 32 (2014), 15–22.
- [55] Miriah Steiger, Timir J Bharucha, Sukrit Venkatagiri, Martin J Riedl, and Matthew Lease. 2021. The psychological well-being of content moderators: the emotional labor of commercial moderation and avenues for improving support. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–14.
- [56] Joseph Turow, Michael X Delli Carpini, Nora A Draper, and Rowan Howard-Williams. 2012. Americans roundly reject tailored political advertising. (2012).
- [57] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *proceedings of the eighth symposium on usable privacy and security*. 1–15.
- [58] Nicholas Vincent, Brent Hecht, and Shilad Sen. 2019. “Data strikes”: evaluating the effectiveness of a new form of collective action against technology companies. In *The World Wide Web Conference*. 1931–1943.
- [59] Nicholas Vincent, Hanlin Li, Nicole Tilly, Stevie Chancellor, and Brent Hecht. 2021. Data leverage: A framework for empowering the public in its relationship with technology companies. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 215–227.
- [60] Richmond Y Wong, Deirdre K Mulligan, Ellen Van Wyk, James Pierce, and John Chuang. 2017. Eliciting values reflections by engaging privacy futures using design workbooks. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 1–26.
- [61] Matthew Wood, Gavin Wood, and Madeline Balaam. 2017. “They’re Just Tixel Pits, Man” Disputing the ‘Reality’ of Virtual Reality Pornography through the Story Completion Method. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 5439–5451.
- [62] Yuxi Wu, Sydney Bice, W Keith Edwards, and Sauvik Das. 2023. The Slow Violence of Surveillance Capitalism: How Online Behavioral Advertising Harms People. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1826–1837.
- [63] Yuxi Wu, W Keith Edwards, and Sauvik Das. 2022. “A Reasonable Thing to Ask For”: Towards a Unified Voice in Privacy Collective Action. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [64] Yuxi Wu, W Keith Edwards, and Sauvik Das. 2022. SoK: Social Cybersecurity. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1863–1879.
- [65] Yaxing Yao, Davide Lo Re, and Yang Wang. 2017. Folk models of online behavioral advertising. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 1957–1969.
- [66] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. 2021. What makes a “bad” ad? user perceptions of problematic online advertising. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–24.
- [67] Yixin Zou, Abraham H Mhaidli, Austin McCall, and Florian Schaub. 2018. “I’ve Got Nothing to Lose”: Consumers’ Risk Perceptions and Protective Actions after the Equifax Data Breach. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS) 2018*. 197–216.

Received January 2024; revised July 2024; accepted October 2024