Privacy Cards for Surfacing Mental Models and Exploring Privacy Concerns: A Case Study of Voice-First Ambient Interfaces with Older Adults

Andrea Cuadra
Olin College
Needham, MA, USA

Samar Sabie
University of Toronto
Mississauga, Ontario, Canada

Yan Shvartzshnaider
York University
Toronto, ON, Canada

Deborah Estrin
Cornell Tech
New York, NY, USA

Abstract

We investigate the ethical and privacy implications of voice-first ambient interfaces (VFAIs) for aging in place through an in-depth engagement with five older adults. Our participants were in the process of becoming experienced VFAI users, and had used a VFAI-based design probe for health data reporting. We create and iteratively refine an interview protocol using *Privacy Cards*. We customize *Privacy Cards* by drawing on participants' previous interviews and device usage logs. Using *Privacy Cards*, we conduct interviews to surface their mental models, and explore their privacy concerns. We find insufficient mental models for proper consent. For example, participants did not know who could access their data, and experienced difficulty distinguishing built-in functionality from third-party apps. Participants initially expressed little worry about VFAI-related ethical concerns, but interviews with *Privacy Cards* revealed nuanced issues, resulting in various implications for future research and design.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**; **User centered design**; **Participatory design**; • **Security and privacy** → **Usability in security and privacy**.

Keywords

Older adults; longitudinal study; smart speaker; voice assistant; Alexa; privacy; ethics; cards

ACM Reference Format:

Andrea Cuadra, Samar Sabie, Yan Shvartzshnaider, and Deborah Estrin. 2026. Privacy Cards for Surfacing Mental Models and Exploring Privacy Concerns: A Case Study of Voice-First Ambient Interfaces with Older Adults. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 19 pages. https://doi.org/10.1145/3772318.3791893

CHI '26, Barcelona, Spain

ACM ISBN 979-8-4007-2278-3/2026/04
https://doi.org/10.1145/3772318.3791893

1 Introduction

As the proportion of older adults continues to increase, so do the technological innovations that promise to support them as they age in place. Voice assistants (e.g., Amazon's Alexa, Google's Assistant, Apple's Siri, ChatGPT voice), which can serve as VFAIs, may have exciting applications for decreasing isolation, motivating physical activity, providing continuity of care, calling for emergency assistance, among others. VFAIs[1] are ubiquitous and pervasive information hubs.

However, VFAIs' inner workings are difficult to discern, and are being increasingly deployed to unknowing users (e.g., older adults in nursing homes or senior centers). VFAIs introduce risk of harm, such as by sharing information across space, groups of people, and time in ways that can threaten privacy and generate other ethical issues. For example, any person's interactions with an Amazon Alexa[2] are by default saved and recorded to the cloud on its associated Amazon account. This means that specific requests made to the VFAI, the tone used to make these requests, any background noises that were audible during the request, and all conversations that ensued during instances when the VFAI may have been activated by accident are accessible to anyone with the password to that Amazon account. A person (e.g., child, grandchild, technical helper) with this password may have access to an older adult's seemingly private interactions captured through their VFAI. Many VFAIs, despite having a single voice, rely on third-party applications from thousands of creators across various contexts, such as health and education, each capable of collecting information about how users interact with their devices from the privacy of their homes. Some of this information is collected without necessarily granting the individuals whose data is at play the ability to delete it, modify it, or know about it.

Moreover, computational models can make far-reaching inferences about people from their data in ways that are difficult for individuals to predict [54]. Relying on history to signal the future, the appearance of increasingly capable technologies will continue to cross, blur, and shift existing contextual boundaries, exacerbating users' vulnerability to privacy violations. As some scholars argue,

[1] VFAIs utilize ambient intelligence [25], which employs advanced sensors and sensor networks, pervasive computing, and artificial intelligence to (passively) integrate into daily human life by making people's surroundings flexible and adaptive. In this work, we use commercially-available smart speakers with screens to study VFAIs.
[2] Amazon Alexa is chosen as the example here due to its extensive use and popularity.

merely asking participants if they have privacy concerns is not sufficient and can be potentially misleading [8, 51], because there is little opportunity to ground responses in participants' lived experiences in an interactive way that is simultaneously reflective, participatory, and theory-informed. Playful and participatory methods that are more reflective and grounded in lived experiences exist, such as the cards that Gaver et al. [35] use as design probes; however, they tend to be more generative, making them unsuitable for eliciting privacy judgments in a theory-informed manner that facilitates clarification and revision. This creates an urgent need to innovate within Human-Computer Interaction (HCI) to design alternative ways of exploring the risks and mitigate potential harms, particularly for groups that have not been well-represented in the design of these technologies.

Responding to this need, we create *Privacy Cards* by taking inspiration from the contextual integrity (CI) theory of privacy [53]. The CI theory defines privacy as the appropriate flow of information, where appropriateness is determined by established contextual information norms (or privacy norms). According to CI, privacy is potentially violated when information flows breach these norms. *Privacy Cards* specifically draw on the CI framework to capture information flows using five essential parameters, allowing for a flexible and systematic way to express information flows and evaluate their alignment with existing privacy norms.

Cuadra et al. [26] discuss the promise of VFAIs for aging in place, describing the stories of five older adult participants who became voice assistant users through their study and with whom they speculated about future interfaces for health and wellbeing through three interviews. The authors caution that VFAIs' pose privacy risks, calling for thoughtful, careful, and systematic consideration of these issues—yet they do not provide findings on privacy itself. With IRB approval for sharing data, we use the interview recordings and device usage logs from Cuadra et al. [26] to inform our study with the same participants. **Our goal is to create and use *Privacy Cards*—a novel, card-based interview protocol—as both a methodological and exploratory tool to surface privacy perceptions and ethical challenges in older adults' interactions with VFAIs, while also advancing inclusive and ethical design practices for emerging ambient technologies.** We chose this approach because existing options were either too complicated for our participants, or not specific enough to address our concerns of potential VFAI harms. We chose design cards because they provide great promise for encoding design knowledge [1], utilizing a human-centered approach to regulation [50], and broadening participation and sharing of power [13, 31] in a simple and user-friendly manner. We specifically investigate the following research questions:

RQ1: Methodological focus. How can the *Privacy Cards* interview protocol help researchers elicit older adults' mental models[3] and concerns about privacy in VFAIs?

RQ2: Ethical focus. What ethical and privacy-related issues emerge when older adults engage with VFAIs using the *Privacy Cards* protocol?

RQ3: Implications. How might insights from using *Privacy Cards* inform the ethical design and study of VFAIs for older adults?

[3]A mental model is a user's understanding of a system [21].

We make three main contributions to CHI: we 1) create *Privacy Cards* and describe their iterative design process, which explains how we adapted them to older adults based on their usage logs; 2) uncover several ethical concerns of using VFAIs for older adults, such as insufficient mental models that affect their ability to properly consent; and 3) identify implications for design and research, including using *Privacy Cards* to inform large-scale survey design. Our novel empirical findings and their associated implications for the research and design of VFAIs provide evidence of the effectiveness of *Privacy Cards*. This exploratory work is an imperative stepping stone towards designing more ethical VFAIs for older adults.

2 Related Work

Our work lies at the intersection of privacy ethics, CI theory, research on older adults' use of VFAIs, data inference, and card-based design methods. While each of these areas has been studied independently, we connect them to address a central challenge: how to co-create meaning and insights with older adults about opaque data practices in everyday interactions with VFAIs using a theoretically grounded and participatory method. We argue that CI is powerful but abstract, older adults are an important yet underrepresented group of users in digital privacy studies about VFAIs, data inference makes risks invisible, and card methods can bridge this—but only if designed considering all of these intersections.

2.1 Existing ethical concerns about digital privacy

A key issue with digital privacy is that we do not yet have ethical, privacy-preserving solutions that adequately address the risks introduced by ubiquitous information technologies. Questions surrounding deception, stereotypes, privacy, and accountability permeate the literature [49, 72, 75]. Moreover, there are well-documented gaps in users' understanding of VFAIs that are already in widespread use [40], increasing the risk of privacy violations. These gaps may be intensified for groups that have been systematically excluded from the design of digital technologies, such as older adults [76]. At the same time, older adults refine mental models through exploration and flexible support strategies, though the design of the technology strongly influences how accessible and effective those strategies are [66].

Despite the increasing interest in exploring privacy design from an HCI perspective [48], these issues remain largely understudied. Several attempts move us closer to better solutions, but the wickedness of the problem [63] requires much more. For example, Alhirabi et al. [3] implemented and evaluated a proof-of-concept prototype to facilitate the incorporation of privacy specific design features into the IoT application from the beginning rather than retrospectively, an important step but certainly not enough.

The information that people disclose to machines, either directly or indirectly, can be used in both good and bad ways. For example, knowing that you may have a certain illness sooner rather than later, can result in better health outcomes. Similarly, knowing a person's interests by analyzing what content or activities they have been engaging with can help tailor future information and recommendations to those interests, making future experiences

more enjoyable. However, these usage patterns can also expose a person's vulnerabilities and fears, which can be nefariously used to many ends, such as fueling panic and division, or tricking people into paying for goods or services they do not want or need. This information can also be problematic if it is used in a harmful manner (e.g., by a predatory insurer). These sorts of assessments can also result in decisions that affect people without their knowledge, and without giving people the ability to revise errors or exercise their rights [24]. Information revealed to a machine can become part of a larger repository of information that is managed by entities with commercial interests, which can be misaligned with the interests of the individual disclosing that personal information [84]. Often, the information shared with the machine is stored, and available to others [28], likely unbeknownst to the user.

Furthermore, some argue that the ability of technology to manipulate information, bypass long-standing privacy values, and influence behavior through social machines cannot be adequately addressed by traditional privacy protection mechanisms [20]. It is unclear which legal regimes should govern these technologies and what consumer protection rules for them should look like [39].

In this work, we respond to these concerns by using a design-focused approach to explore, identify, and unveil potential harms in a highly contextualized manner with older adult participants, generating rich questions for further exploring the ethics of VFAIs.

2.2 Contextual integrity (CI) theory of privacy

Privacy concerns are not easily measurable, calling for more nuanced treatments of the notion of privacy within HCI [8]. One way to achieve this is by drawing from theoretical frameworks, such as CI [53]. CI can help researchers explore privacy norms in social contexts. Privacy norms shape our expectations for what is appropriate in a given situation—they are socially constructed and dynamic, and they may change when new technologies create new social contexts [61].

CI has been used in HCI for different purposes. For example, Shvartzshnaider et al. [67] used a CI-based vignette study to understand privacy norms in an educational context. Apthorpe et al. [4] used a similar approach to study privacy expectations of Smart Home devices. In further work, Apthorpe et al. [5] used the CI-vignette study method to understand parents' privacy expectations about the use of devices protected by the Children's Online Privacy Protection Rule. To explore the ethical dimension of big data research project, Zimmer [83] developed a CI-based heuristic to help researcher think about relation of privacy, anonymity and harm.

We chose to create *Privacy Cards*, because there are not many easily accessible strategies to facilitate the use of CI for developing a nuanced understanding of the *"why"* for groups that have been marginalized and underrepresented in the design of digital technology. In today's world, avoiding using digital technology is increasingly difficult, making it more urgent to uncover privacy norms in different demographics and cultural contexts. Recently, Kumar et al. [48] developed a roadmap for HCI researchers on using CI in conducting qualitative privacy research. This work serves as an example of this roadmap in action towards guiding HCI and social computing researchers on how to apply CI to qualitative projects.

2.3 Older adults' use and perception of VFAIs

Although work examining older adults' use and perceptions of VFAIs has increased in recent years [17, 23, 26, 38, 44, 52, 58, 59, 69], older adults remain substantially underrepresented in VFAI research. Reviews of the computing literature found that 97% of ACM Digital Library publications did not consider age or aging in voice user interface studies [73, 74]. At the time these reviews were conducted (2019–2020), nine percent of the global population was over 65[4] , indicating roughly a threefold underrepresentation. As a result, VFAIs do not adequately represent the needs of older adults, which may put them at higher risk of harm.

Stigall and Caine [73] further observed that most studies reported older adults taking more time to operate voice user interfaces and/or making more errors than younger adults, with a minority finding no age-related differences. A more recent systematic review highlights older adults' visions for future AI-based conversational systems, expressing a desire for more human-like interactions, personalization, and greater control over their information [42]. The desire for more human-like interactions and personalization are notable, because they may increase the amount of trust in the VFAIs alongside the amount of data collected. Moreover, designing for greater control over user's information requires detailed understanding of their conceptual models of the technology. Our work dives deeper into these matters while especially considering the emotional limitations and potential dangers associated with anthropomorphized conversational agents [27, 29].

Other studies document a mismatch between media portrayals of VFAIs and older adults' lived experiences. Sin et al. [68] found that older adults perceived VFAIs as more primitive, limited, and difficult to set up than mass-media narratives suggest. Such discrepancies may create feelings of inadequacy and reinforce stereotype threats around older adults' technology use. Relatedly, Shandilya and Fan [65] found that while older adults expressed enthusiasm for AI-enabled products, including VFAIs, they also worried about privacy intrusions and a loss of agency.

Research specifically examining older adults' privacy concerns with VFAIs remains relatively scarce. A comparative study of adults under 65 and 65+ found that older adults reported fewer VFAI-related privacy concerns regarding consent, data security, and data protection [71], raising questions about whether they fully understood the associated risks—a theme we investigate here. Similarly, Bonilla and Martin-Hammond [12] found that older adults expressed concerns about VFAI data practices but were often unaware of available privacy resources, concluding with a call to better support this population. Finally, So et al. [69] found that older adults had many privacy and ethical concerns about future VFAI technologies, but did not explore their mental models about how VFAIs currently work. Our work addresses these concerns in a novel manner by examining how *Privacy Cards* can surface gaps in understanding and reveal the ethical challenges older adults face when interacting with VFAIs.

2.4 Data inference

The ability to produce inferences from data elevates the need to approach this problem in theory-grounded, participatory manner.

[4]https://data.worldbank.org/indicator/SP.POP.65UP.TO.ZS

Health status information may be inferred from inconspicuous sensors embedded in our built environments [57], wearable technologies [33, 43], digitally mediated social networks [45], keystroke patterns [79], and various other not necessarily medical contexts. According to Turow [77], many scientists believe that a person's weight, height, age, and race, potential illnesses they may have, can also be identified from the sound of that individual's voice. Moreover, human machine interactions and information disclosures go beyond the individual, meaning that data a person discloses can inadvertently disclose information about others as well [22, 54, 78]. VFAIs in particular contain a large amount of information about people, such as the particular times of day when people use certain commands, users' explorations (or lack thereof) of new topics, and how their usage patterns relate to those of other users [10].

The amount of information we can derive from inconspicuous, or seemingly innocuous data sources, highlights the importance of studying potential harms of emerging technologies, in particular for marginalized groups, such as older adults. This work illustrates a way to do so, and can serve as an initial foundation to determine how to mitigate potential harms.

2.5 Use of card activities to inform digital design

HCI has long concerned itself with the use of cards as a design tool, especially as a way for designers to encode and communicate design knowledge [1]. Over time, card-based approaches have expanded to support a wide range of design goals. Tools such as *Envisioning Cards* [34] and the *IDEO Method Cards*[5] illustrate how cards can catalyze creativity, surface values, and prompt new perspectives. Synthesizing this landscape, Hsieh et al. [41] identified seven forms of design knowledge embedded in 161 card decks—from creative inspiration to values in practice—and noted that most decks emphasize early-stage ideation. Our *Privacy Cards* may support design at all stages. Card sets have also been used to support ethical and inclusive design, such as *Playing the Legal Card* [50] and *Responsible & Inclusive Cards* [31], which help practitioners engage with challenging topics and democratize participation. Other HCI decks illustrate this breadth: the *Tarot Cards of Tech* [6] use speculative prompts to help teams anticipate unintended consequences; *Judgment Call The Game* [7] engages players in scenario-based ethical decision-making; and the *Building Utopia* toolkit [16] supports community-led envisioning of equitable technological futures. Together, these examples highlight cards' flexibility as reflective, ethical, and participatory design tools.

Early work such as Gaver et al. [35] introduced Cultural Probes—open-ended postcard kits intended to spark inspiration and surface unexpected ideas with older adults. While powerful for generative design, their ambiguity makes them unsuitable for surfacing or structuring the nuanced, theory-driven privacy judgments required in our work. More recent efforts explore card-based methods for eliciting such judgments. For example, Berkholz et al. [11] created *Privacy Taboo*, a card game designed to serve as a playful breaching interview method, fostering discourse on unwritten privacy rules. When playing *Privacy Taboo*, their younger adult participants articulated their information needs when consenting to fictive data requests, even when contextual cues were limited. Our *Privacy Cards* builds on these traditions by grounding elicitation in older adults' lived experiences while also systematically leveraging a privacy framework to surface concrete information needs. We rely on cards' playful design affordances to foster engagement and make discussions of challenging topics more accessible, drawing directly from participants' VFAI usage logs and interviews. By doing so, we tie together privacy ethics, the CI theory, research on older adults' use of VFAIs, and data inference risks towards creating a deeper understanding of ethical concerns that regard older adult use of VFAIs.

Taken together, prior work highlights a critical gap at the intersection of the domains described in this section. CI offers a rigorous framework for reasoning about privacy, and card-based methods provide accessible, participatory tools for ethical engagement. Yet CI has, to the best of our knowledge, not been operationalized in a way that is specifically tailored to older adults' interactions with digital technologies (less so with VFAIs specifically). Meanwhile, studies of older adults' VFAI use document privacy concerns and desires for control, but lack methods for uncovering the conceptual foundations required for a meaningful conception of privacy. The need for such a method is amplified in light of data inference practices that make information flows invisible. In this work, we use the CI to create *Privacy Cards*, which addresses this critical gap by building on prior work to surface older adults' mental models and privacy concerns in real contexts.

3 *Privacy Cards* design process

To design *Privacy Cards* we combined user tests, senior center field observations, and consultations with experts in CI, older adults, and individuals outside computing or adjacent fields. Additionally, we incorporated information from earlier participant interviews and usage log tracking.

This study relies on other research about designing VFAIs to support aging in place that was led by the same first author, where the authors conducted a two-month field deployment of a VFAI [26]. During this period, they tracked participants' usage of their VFAIs and interviewed them on different topics. Their interviews covered the following topics in chronological order: 1) general use and familiarization, 2) home health, and 3) wellbeing. Cuadra et al. [26] shared their primary data with us with permission from all IRBs involved. In this paper, we report the findings from a privacy interview that we carried out between their home health and wellbeing interview, using familiarity from the prior engagements and device usage logs.

We now describe our design process by topic. First we briefly explain CI and the five parameters that it requires (Sender, Recipient, Subject, Attribute, Transmission Principle). These parameters are the core of CI, so they became the core of *Privacy Cards*. We describe the details of how we adapted the terminology used in CI to make it more accessible to a wider audience. For example, from talking to older adults and individuals outside computing or adjacent fields, we determined that CI uses academic language that can be difficult for many to understand, so we worked to simplify the language. We also selected values for each parameter that we considered would be relevant to our participants' own interactions with their VFAIs.

[5] https://www.ideo.com/journal/method-cards

Finally, we describe how we developed the mechanics of *Privacy Cards* and how we then digitized them.

3.1 Contextual Integrity

We initially studied Nissenbaum's seminal book, *Privacy as Contextual Integrity* [53]. CI posits that privacy is not maintained by keeping an information secret, adhering to a well-defined procedure, providing specific access controls, or seeking informed consent. CI defines privacy in terms of the appropriateness of information flows in a given context, as prescribed by governing contextual norms. Potential privacy violations occur when information flows deviate from established norms or societal expectations. As connected technologies are introduced into our homes, our information is exposed beyond the home context, connecting a typically personal indoors setting into a globally reachable and observed environment. This notion of the "appropriateness of an information flow" is ever more important when technologies generate information flows that users are unaware of and that may violate the privacy norms of a specific context.

To examine the privacy implications of these technologies, the CI framework identifies information flows that deviate from the norms and uses the CI heuristic to evaluate their contribution to the purpose, values, and function of a given context. CI analysis requires capturing individual information flows and norms using five parameters:

(1) **Sender:** the actor that sends the information
(2) **Recipient:** the actor that receives the information
(3) **Subject:** the actor of whom the information is about
(4) **Attribute:** this is the information type, in our study we call it "data"
(5) **Transmission Principle:** this states the constraints on the flow, such as sharing confidentially or with subject's consent.

A noteworthy aspect we debated was how to categorize the VFAI. We ultimately decided to categorize it as a communication medium, like a telephone device.

3.2 Adapting CI to cards for older adult participants

We began by mapping each of the five CI parameters (Sender, Recipient, Subject, Attribute, Transmission Principle) into a card category. In other words, each card category was determined by a CI parameter. We then simplified the language to be more accessible to our older adult participants. For example, each card had a CI parameter with an associated value (e.g., "Your doctor", or "You") expressed in simplified language: "Who Sends" instead "sender", "Who Receives" instead of "recipient", "Data" instead of "attribute", and "Condition" instead of "transmission principle." The CI "subject" parameter was implicitly the participant, since questions were posed using the second-person pronoun "you," see Figure 2. The use of "you" as a subject helps elicit participants' individual preferences, which could point to the gap between those preferences and what is normatively acceptable.

We made some changes based on data from the first interviews conducted by Cuadra et al. [26]. First, we adapted the language to use the words the participants used during these interviews. For example, one participant would refer to activating the VFAI as "calling" the VFAI, so we wrote on the Data card for that attribute "The timestamps of when you've called [the VFAI]" instead of "The times you have used your [VFAI]." Second, we incorporated participant usage logs to account for their specific experiences with the VFAIs by using the names of third party voices they had been using for some Who Receives cards: "Voice Apps, LLC. the creators of Sleep Sounds: Thunderstorm Sounds" (P1); "Spotify, the company that sends songs to your [VFAI]" (P2, P5), "Avocado Labs, the creators of: 6-Minute Full Body Stretch" (P3), and "Matchbox.io Inc, the creators of: Daily Stretch" (P4). We also tailored our cards to reflect participants' prior experiences in the Cuadra et al. [26] study. For instance, a design probe for health data reporting, which asked about activities of daily living, was represented as a Data card labeled "responses to a health-related questionnaire your doctor requested" (see Figure 2).

3.3 Determining CI parameter values

One of the key design challenges we faced was determining the content, or the values, for each card category/parameter. For example, in our first version, which was informed from visiting a senior center with a program to provide VFAIs to some of its members, we imagined having multiple *senders*, including: the VFAI[6] , the senior center tech support person, the senior center director, the participant's doctor, child, paid home health aide, and friend. We additionally had a list of 12 *recipients*, which included the senders in addition to the the local police, government intelligence agencies, an Amazon employee, the participant's connections on social media, and a social worker. We also had 11 *attributes* (e.g., "the times and dates when you have asked for the weather", "a list of people you have called using [the VFAI]") and 11 transmission principles (e.g., "if you are notified", "if the information is not stored", "in an emergency situation"). This list underwent several iterations, because we needed to narrow down the number of information flows to fit in an hour-long interview with our older adult participants. The main criteria we used to narrow down the list was 1) closeness to participants' interactions with the VFAI in our study, and 2) relatedness to the research topics we were interested in (e.g., VFAIs for home health). In the second version, we significantly trimmed down the number of questions, prioritizing scenarios we were specifically building in our lab, including a voice app for health data reporting. The final deck we used for this study, described in Section 4.2.2, had one *sender* (always attached to the VFAI *transmission principle*), four *recipients*, four *attributes*, and 18 optional *transmission principles*. Note, the values of the parameters used, not the parameters themselves, can vary depending on a researchers' goals.

3.4 Developing the mechanics for *Privacy Cards*

We made several design decisions to develop the card sorting mechanics. For example, we created some orientation activities to establish conversational grounding with the participants, including: 1) showing how to move a card, 2) asking about each person/role in the Who Receives cards, and 3) ranking the contextual relevance of a card by placing a Data card on a scale from "this information

[6]This would be inconsistent with the CI theory, which requires the actor values expressed in roles and capacities.

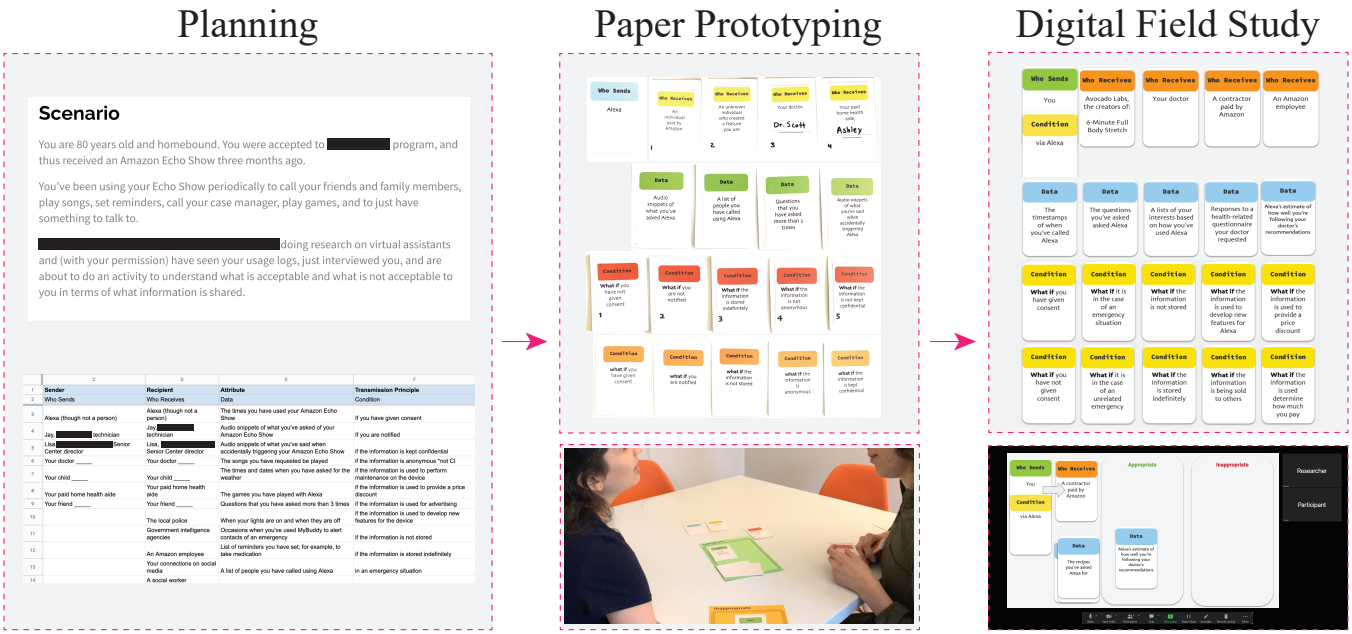

Figure 1: Note, this image is an embedded PDF, so zooming in may help see the graphics at a more appropriate size. This diagram depicts three phases of the *Privacy Cards*' design process: 1) Planning—Hypothetical scenario as a prompt for getting feedback from CI experts (top), and CI parameters and their respective values in the context of older adults and VFAIs at a senior center (bottom), 2) Paper prototyping—Cards generated from the planning phase after narrowing down the sender options and creating two sets of conditions (what is expected to lead to appropriate, or inappropriate evaluations), and 3) Digital Field Study—A portion of the final set of *Privacy Cards* we used over Zoom with older adult participants in their homes.

is irrelevant to this person" to "this information is relevant to this person" by *recipient* (see Section 4.2.2 for more details).

Another example is that we were unsure whether it was best to cover breadth (i.e. all available *recipients*) or depth (i.e. all available *transmission principles*). We ultimately did a combination of both, including at least two *recipients* (Amazon employee or contractor, and doctor), and the *transmission principles* that naturally emerged during the conversations that ensued during the activity. We still had all the options available, and in some cases, such as when a participant chose to interview over the agreed upon hour, we were able to cover more content.

To further narrow down the number of flows and meet our study's time constraints, we fixed some parameter values. The *subject* and *sender* were always the participant, and every information flow included a *transmission principle* stating that the data was sent via each participant's VFAI. We also carefully arranged the cards and sorting activities based on the information flows we prioritized for the study, and in the order we perceived would better suit our participants' card sorting experience.

Finally, we added background buckets to guide the movement of the cards. As depicted in Figure 2, these background buckets had labels corresponding to the participant's evaluation of an information flow, such as "appropriate" or "inappropriate". For the step depicted in Figure 2, the stack of blue Data cards was originally placed near the bottom edge of the screen under the orange Who Receives cards, and participants could move them to either the "appropriate" or "inappropriate" bucket, or somewhere in between. Given that *transmission principles* impose constraints on the flow, we had a second stage setup where participants could re-evaluate Data they had already considered appropriate or inappropriate based on *transmission principles*, or Condition cards, as portrayed in Figure 3. In this case, the yellow Condition cards are stacked on the top right of the screen.

3.5 Digitizing *Privacy Cards*

Even though we initially developed the cards to be used in-person (see Paper Prototyping Phase in Figure 1), we had to transfer them to a digital format so we could run the study over Zoom during the COVID-19 pandemic. To do so, we recreated the cards on Adobe Illustrator and exported them as PNGs, which we then uploaded to Google Slides. We shared the part of the screen with the slide, and moved the cards around based on participants instructions in real-time via Zoom screen-sharing. The flows were verbally described for one participant who did not have access to Zoom for the visual elements of the activity. Even though having to digitize the cards presented some challenges, it ended up improving *Privacy Cards*. Unlike the original, paper-based cards, the digital ones allowed us to reach participants remotely. This also presents the possibility of using protocols similar to *Privacy Cards* at larger scales.

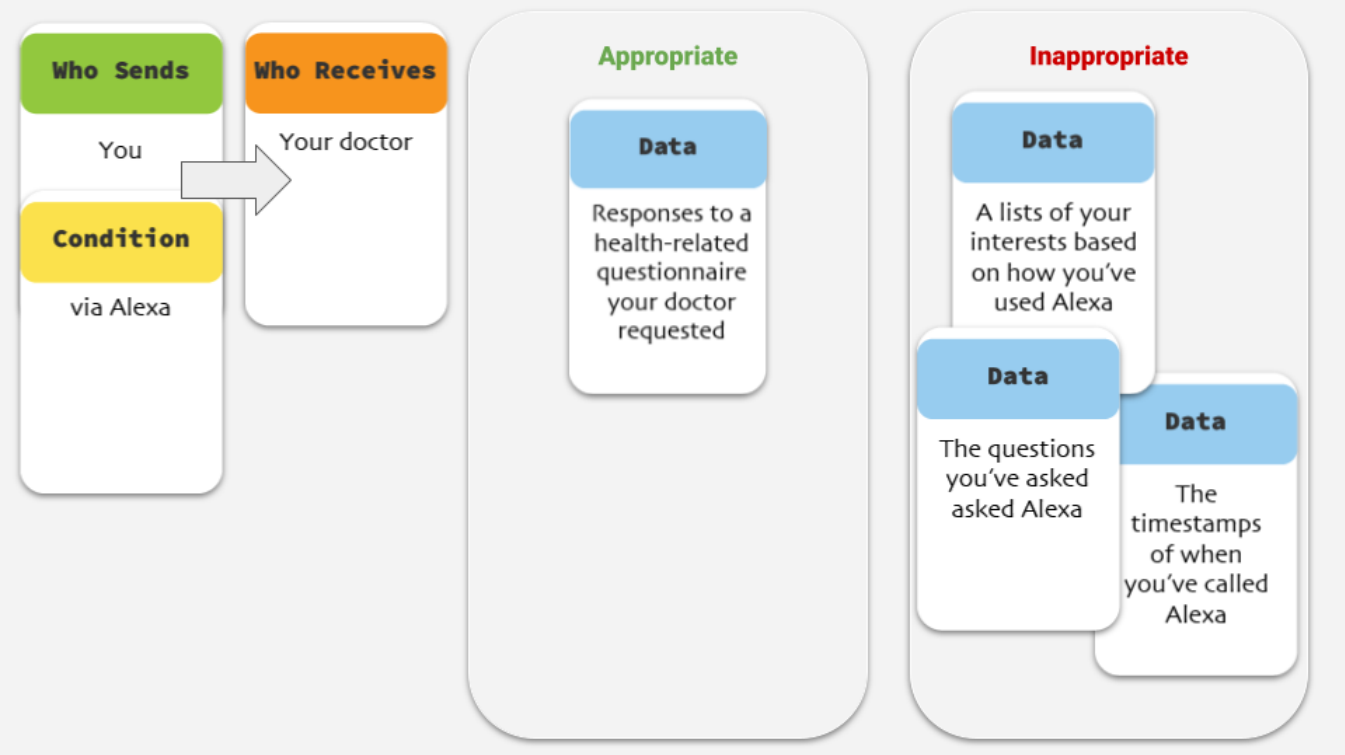

Figure 2: *Privacy Cards* in action during step four of P3's session. Participants used the cards to evaluate the appropriateness of different information flows. The VFAI in the cards was labeled by the name the participant used to refer to it (i.e. Echo or Alexa).

4 Field Study Method

We conducted an IRB-approved (Cornell University, #1912009271) exploratory field study with older adults living alone (N=5), recruited via local senior centers in New York City. The study was conducted remotely via phone or Zoom interviews, all conducted by the first author. We now describe our participants, procedure, and analysis approach.

4.1 Participants

Five older adults (four women, one man) between the ages of 62 and 85, with varying degrees of technological expertise, were recruited. Recruitment focused on people who belonged to older adult communities [14, 62] via senior centers. Senior centers are community centers designed to make older adults feel supported and happy—they bring older adults together for a variety of services and activities designed to enhance their quality of life [9]. A short presentation about the study was given during a Zoom meeting with many senior center directors, and each director was sent a flyer with details about the study to share with their members. Some directors responded with the names and phone numbers of prospective participants. Each prospective participant was then called, provided with details of the study, and given the opportunity to ask questions. They expected the call because their senior center director had informed them in advance. If a person expressed interest in participating, a time was arranged to drop off the VFAI and obtain consent.

This study responds to Dix [30] who argues for the value of small-scale studies "as we move from a small number of applications used by many people to a 'long tail' where large numbers of applications are used by small numbers of people," and Vines et al. [80], who suggest critical engagement with an individual's context as a strategy to combat common stereotypes that prevail in the literature. Limiting the number of participants to five allowed for the creation of deep, personalized engagements as described by Cuadra et al. [26].

Building on prior literature, the study also aimed to focus on older adults who have been historically underrepresented in the design of technology [37]. Gell et al. [36] found higher prevalence of technology use in older adults with five characteristics: younger age, male sex, white race, higher education level, and being married (all p values <.001). Thus, participants were recruited who did not have more than two of the five characteristics associated with higher prevalence of technology use—none of the participants were married, the younger ones (in their sixties) were not white, and the only man was low-literate. All participants lived independently, by themselves, and had WiFi in their homes. See Table 1 for demographic details by participant.

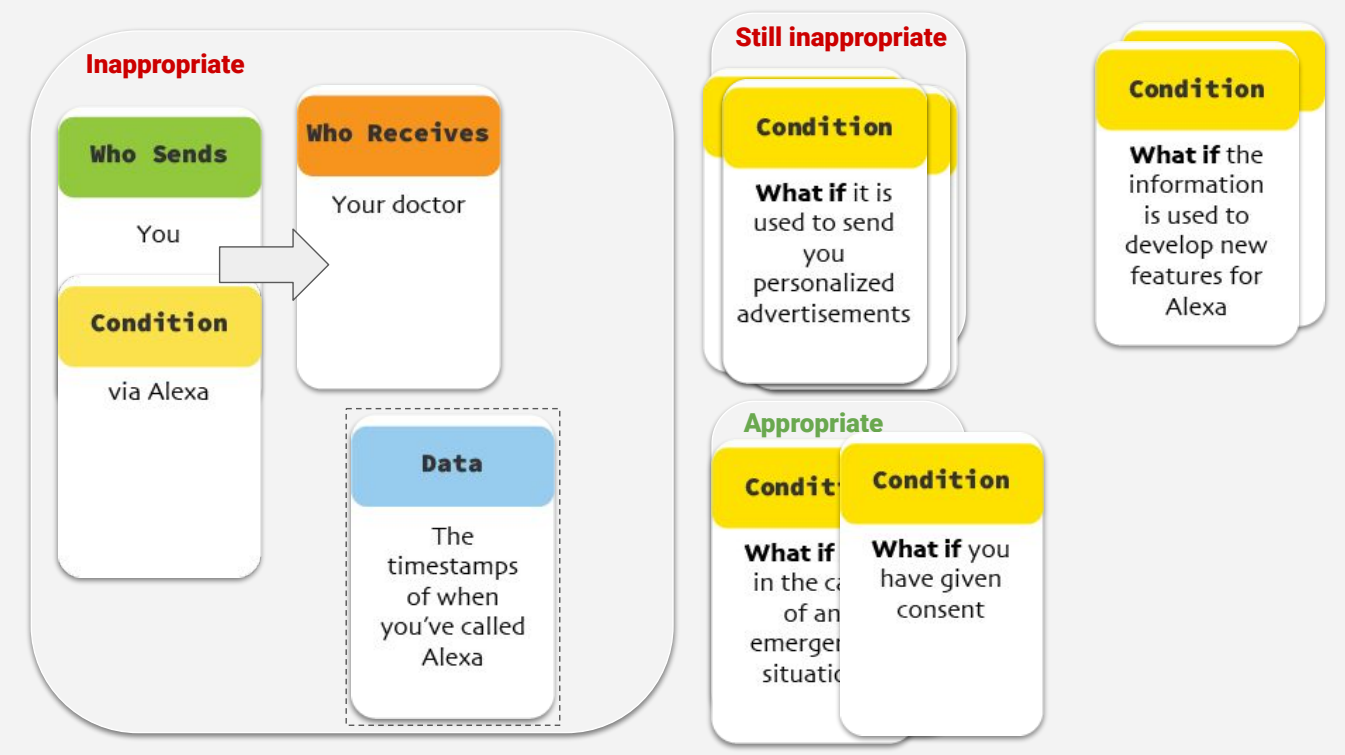

Figure 3: *Privacy Cards* in action during step five of P3's session. After their first appropriateness evaluation, participants were asked to reconsider them based on new "conditions" imposed on the flows. We placed a Data card that was considered inappropriate to complete a four-card flow on the left side, and had Condition cards stacked on the top right of the screen. The flow is then re-evaluated considering these conditions.

Table 1: Summary of participants, their engagement with their VFAIs, and their interview durations.

P #	Age	Gender	Race	Median # of weekly activity (first 5 weeks)	Days with VFAI since 1st activity	Interview duration
P1	67	M	Black	71	26	52m
P2	82	F	White	36	11	1hr 13m
P3	85	F	White	37	43	2hr 10m
P4	85	F	White	45	9	1hr 15m
P5	62	F	Latinx	41	118	1hr 05m

4.2 Procedure and materials

We now describe the field study's procedure, cards and mechanics, interview guide, hardware and software, and data collection.

4.2.1 Procedure overview. A VFAI was dropped off at each participant's home, and participants were asked to interact freely with it. They were given training on common usages (e.g., the weather, music, and information retrieval), were shown how to mute the device if they did not want it to be "listening," and had any questions answered. Participants interacted with the VFAIs for at least five weeks and engaged in three touchpoints with a researcher before the *Privacy Cards* interview: 1) drop-off, 2) a familiarization period interview that was unbiased by any design interventions, and 3) an interview based on interactions with a home health voice app design probe. The interviews leading up to the interview session are reported in [26] (see Table 2).

The *Privacy Cards* sessions were conducted either via phone calls or via Zoom video conferencing meetings (depending on the participant's preferences and abilities). The shortest interview was 52 minutes and the longest one two hours and ten minutes. For each session, we used a slide deck with *Privacy Cards* to discuss privacy grounded in the context of participants' interactions with the VFAIs. For example, a question we asked all participants through this activity was, "would it be appropriate or inappropriate for you

Table 2: Summary of interviews.

Familiarization (reported in [26])	Home Health (reported in [26])	*Privacy Cards*
Alexa strengths and challenges	General use	General use
Questions from usage logs	Home health voice app	Homework debrief
Brainstorm potential uses	Homework: try voice app once on their own	Privacy concerns
		Privacy Cards activity

(the *sender*) to send your doctor (the *recipient*), via [the VFAI] (a *transmission principle*), your (participant is the *subject*) responses to the health assessment you took last time we talked (the *attribute*)?" Depending on how much time was available, different *transmission principles* (such as "what if the information is being sold to others", or "what if the information is kept confidential") would also be added to the information flows at hand. Once a flow was deemed appropriate or inappropriate, it was placed in its respective bucket. The visualization of these flows (see Figure 2) prompted participants (except for P1, who did not have access to screen-sharing) to reevaluate previous decisions and the researcher to ask more questions about the participants' rationales.

4.2.2 Privacy Cards *used and mechanics.*

Table 3: Number of *Privacy Cards* by category as used in our field study.

Type of Card	Number of Cards
Who Sends	1
Who Receives	4
Data	4
Conditions if Appropriate	9
Conditions if Inappropriate	9

For our field study, each deck had 27 cards (as shown in Table 3). We used *Privacy Cards* following the steps below. Our Supplementary Material includes the slide deck for these steps, and an Adobe Illustrator editable file for the cards. Note that we skipped some of these steps for some participants, such as step three, if the participant was taking too long to complete the activity:

(1) **Slide 1:** Moved one Who Sends (You) card to a box that said "move card here" (see Figure 4).

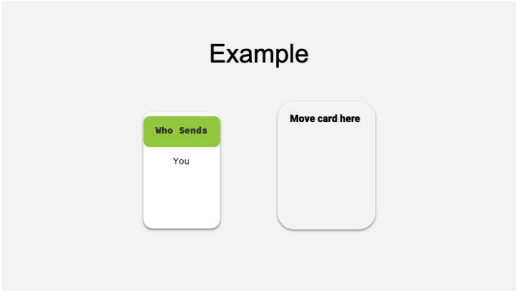

Figure 4: Slide 1: Learning how to move a card.

(2) **Slide 2:** Introduced four Who Receives cards ([specific third-party voice app], your doctor, a contractor paid by Amazon, An Amazon employee), and asked participants to tell us about each recipient and their role to characterize mental models (See Figure 5).

Figure 5: Slide 2: Conversational grounding for *recipients*.

(3) **Slides 3–6:** Asked participants to rank the relevance of four Data cards (The timestamps of when you've called [the VFAI], The questions you've asked [the VFAI], A list of your interests based on how you've used [the VFAI], and Responses to a health-related questionnaire your doctor requested) to the Who Receives cards, from less relevant to more relevant. This was used as a manipulation check to make sure that the specific data types we were evaluating were considered important for the contexts at hand (Figures 6a–6d).

(4) **Slides 7–10:** Changed the slide background to a new setup with three cards fixed on the left side (Who Sends, Who Receives, and Condition: via [the VFAI]). Asked participants to place four data cards in either the "Appropriate" or "Inappropriate" buckets. Time permitting, participants placed Data cards in these buckets for each of the four Who Receives cards (Figures 7a–7d).

(5) **Slides 11 and 12:** Selected an information flow with four cards deemed appropriate or inappropriate (e.g., Who Sends: You, Who Receives: Your doctor, Condition: via [the VFAI], Data: Responses to a health-related questionnaire your doctor requested), and introduced more Condition cards. Decided whether the flows remained appropriate or changed to inappropriate based on the new transmission principles introduced by placing the cards in the "Still appropriate" and "Inappropriate" buckets (see Figure 8a). Finally, the reverse happened but with "Appropriate" or "Inappropriate" buckets flipped (see Figure 8b), and the Data card changed to

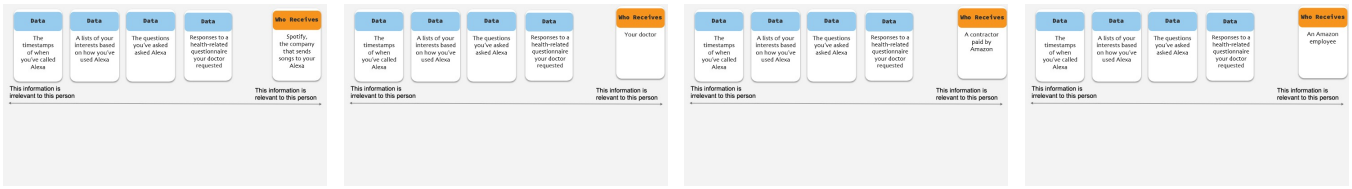

(a) Slide 3: Sorting cards by relevance to third-party company. **(b) Slide 4: Sorting Data cards by relevance to doctor.** **(c) Slide 5: Sorting by relevance to contractor paid by Amazon.** **(d) Slide 6: Sorting Data cards by relevance to Amazon employee.**

Figure 6: Setup slides for step three.

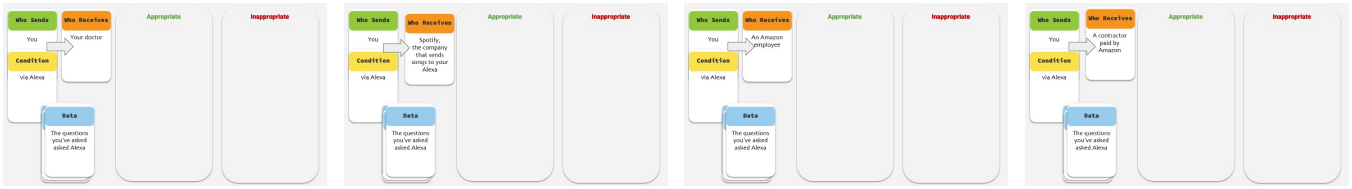

(a) Slide 7: Sorting Data cards sent to doctor. **(b) Slide 8: Sorting Data cards sent to third-party company.** **(c) Slide 9: Sorting Data cards sent to an Amazon employee.** **(d) Slide 10: Sorting Data sent to contractor paid by Amazon.**

Figure 7: Setup slides for step four.

something the participant had originally deemed inappropriate (e.g., Data: The timestamps of when you've called [the VFAI]).

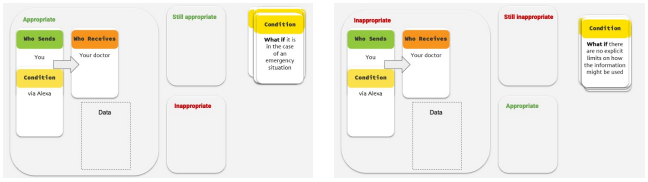

(a) Setup for sorting cards that were deemed appropriate. **(b) Setup for sorting cards that were deemed inappropriate.**

Figure 8: Setup slides for step five.

4.2.3 Interview guide. The interview guide had five main sections: 1) warm-up/rapport building (five minutes), 2) questions about using home health design voice app on their own (five minutes), 3) general privacy concerns (15 minutes), 4) interactions with our tool (30 minutes), and 5) closing remarks (five minutes). For the general privacy questions, we asked questions such as, "Does anything worry you about having a [VFAI] in your home?", "Where is information stored?", and "Who gets to see this information?" Then we started the activity by saying, "Now we are going to do an activity surrounding your opinions about what type of information sharing is appropriate or inappropriate."

4.2.4 Hardware and software. Amazon Echo Shows (2nd Gen), which have a 10.1" HD smart display with Alexa, were selected as the VFAI. Dedicated email and Amazon accounts were created for each device. The default music player was configured to be Spotify, because the Spotify music voice app provided more flexibility than the default Amazon Music. The dedicated accounts allowed usage to be monitored and voice apps (including the health data reporting voice app) to be installed on the devices at any time. The devices were then dropped off at participants' homes and connected to their WiFi.

4.3 Data analysis

Interview transcripts were anonymized and uploaded to a shared spreadsheet document. Each transcript was open-coded [46] for thematic analysis [15] by at least two researchers, who met to discuss how the data addressed our research questions, systematically analyzed participant responses, and grouped them into themes. For example, if a transmission principle emerged during the interview, that was coded as "TP specified." Examples of other codes were "not concerned about privacy," "low tech familiarity," and "the-VFAI-is-great attitude." We also annotated each participant's evaluations of specific information flows based on the transcripts, and backed by the resulting card configuration in Google Slides. The first author reviewed all the coding to ensure consistency and met with the other coders to resolve disagreements. Ultimately, the main themes we identified included: application of CI, flow appropriateness, privacy concerns, insufficient mental models, distinction between different entities at play, emotional attachment and cognitive demand. The data recorded included usage log entries from the research accounts on participants' devices, the total number and general usage are described in [26]. Table 1 shows the median weekly activity or usage logs per participant.

5 Findings

This study highlighted nuances in contextual boundaries and prompted a reassessment of privacy concerns. In this section, we first address

RQ1, which was methodologically focused. We describe how we operationalized CI by using *Privacy Cards*, making otherwise obscure information flows comprehensible to our participants. Second, we share ethical issues that we uncovered, addressing RQ2. We found that limited understanding of the VFAI's operation hindered participants' agency over information flows, which was exacerbated by a lack of awareness regarding the VFAI's ties to Amazon. Throughout this section, we highlight the specific mental models conveyed by our participants and explain how they were insufficient, suggesting both *Privacy Cards*' effectiveness and a problem definition that makes corrective action possible. There is some overlap in the findings between research questions, but we organized them in by placing them where the fit was strongest. We address RQ3 (implications) in the Discussion.

5.1 RQ1: Methodological findings

5.1.1 Participants' evaluations of information flows. Some evaluations of information flows as either appropriate or inappropriate were universally agreed upon by all participants while others had much more variation (see Table 4). All participants thought it was appropriate to send responses to a health-related questionnaire to a doctor but never to an Amazon representative. Information flows such as sharing usage log timestamps, or the questions participants had asked the VFAI, had more variation in how appropriate they were perceived. For example, while evaluating a Data card involving "Matchbox.io Inc, the creators of: Daily Stretch," a voice app she used during the study, P4 asserted, "*[Matchbox doesn't] need to know what my mind is doing: what I want to say, or what I want to learn. All they need is for me to do the exercise.*" Unlike several other participants, P3 did not consider sharing the information in any of the Data cards, other than responses to the health-related questionnaire to her doctor, appropriate. P3 explained that her doctor should not receive information that her doctor did not need to know, "*That's too personal. My doctor doesn't have to know everything about me. Really, what if I'm on a dating site and I don't want him to know? Seriously, I mean, there's a lot of personal stuff that I might be interested in as a user and I don't want my doctor to know that.*" These findings highlight the importance of clearly stating the contextual parameters when describing and evaluating information flows.

5.1.2 Condition cards uncovered values affecting information flow appropriateness and transmission principles. Privacy Cards elicited qualitative responses that uncovered participants' values and areas of uncertainty. As indicated by the asterisks in Table 4, participants were frequently uncertain about the appropriateness of a flow. As the interviewer introduced new cards, or asked more questions related to a particular information flow, participants reconsidered their assessments. For example, P4 initially considered it inappropriate for Matchbox.io to receive information about when she was doing her exercises, but she changed her mind when a new Condition card (or transmission principle) prompted her to consider that having that information may help Matchbox.io to create more exercises. Similarly, P1 initially thought it would be inappropriate to send Amazon the timestamps of when he had used the VFAI. However, when asked *"what if they told you that it did matter to them and that it helps them make a better product if they knew that?"* He changed his evaluation to appropriate, and added a need to be able to consent, *"so they give me a chance to say yes or no, to send it or don't send it."* Again, these findings align with CI in highlighting the importance of stating the values of all required parameters. They especially reveal transmission principles' important to our older adult participants, such as those that could improve the products they used.

5.1.3 The playful mechanics allowed us to identify surprising perceptions. Conversations during interactions with *Privacy Cards* helped unveil values to inform design choices. For example, P2 did not want to inconvenience her doctor, valuing consideration of others' time. She thought that sending information to her doctor via the VFAI would mean her doctor would have to listen to the VFAI, which she deemed inappropriate. However, if the same information were to be conveyed using written form, then the same information flow was deemed appropriate. P2 perceived that would be less intrusive.

Mental model: one to one mapping between modalities

The participant mental model was that there is a one to one mapping between the modality used to interact with the VFAI (speech) and the modality used to consume information from these interactions. However, there is a possibility of information being transformed across modalities—such as speech being converted into text and text into speech during processing and storage. In this speculative example, the spoken responses to the medical questionnaire could result in a text-based PDF file matching existing medical records.

As these examples illustrate, the evaluations were often non-binary, dynamic, and value-based. The Condition cards further affected the complicated nature of these assessments, often swaying them in opposite directions. Furthermore, though systematic, the activity gave room for unstructured or unplanned exchanges, allowing participants to impose additional conditions themselves and provide additional rationales, shedding light into what mattered to them.

5.1.4 Using Privacy Cards *elicited privacy concerns from participants who initially said they did not have any.* Examining perceptions around information flows facilitated via the VFAI using *Privacy Cards* revealed discrepancies in how participants generally thought about their perceptions towards privacy and the VFAI itself. For example, P3 expressed, "*Look, I'm at a stage where I don't care. You want to go find out what I'm saying or doing, it's okay with me. I don't care.*" Despite saying she did not care, when using *Privacy Cards* she expressed significant concern about certain information flows not happening, at least not without her explicit consent. Specifically, P3 thought it was inappropriate to send the timestamps of when she responded to the health questionnaire to her doctors via the VFAI.

Table 4: Appropriate information flows by data type and recipient (all via Alexa). Each entry in the table lists the participants that found this flow appropriate. E.g., sending responses to a health-related questionnaire their doctor requested to an Amazon employee or contractor was considered inappropriate by all participants. An asterisk denotes uncertainty about the appropriateness of the flow, or the specification of certain constrains on the flow.

Data Type \ Recipient	Amazon employee or contractor	Third party company or voice app	User's doctor
Responses to a health-related questionnaire your doctor requested	None	None (P1 & P5 N/A)	All participants
The questions you've asked Alexa	P1, P2, P4, and P5	P4* (P1 & P5 N/A)	P4*
A list of your interests based on how you've used Alexa	P1, P2, P4, and P5*	None (P1 & P5 N/A)	None
The timestamps of when you've called Alexa.	P1*, and P5	P2, P3*, and P4 (P1 & P5 N/A)	P1*, P2*, P4, and P5

Mental model: survey responses submitted via the VFAI do not contain timestamps

The participant applied the governing norms of paper-based forms—where precise completion times are not typically recorded—to a digital questionnaire, where timestamps are automatically and routinely captured. This clash between "paper form" and "digital form" norms led to an unanticipated expansion of information collection and illustrates a deviation in CI.

Similarly, despite also expressing having nothing to hide, P5 deemed many flows of information inappropriate (see Table 4). P5 also expressed not wanting to appear "*crazy*," presumably by friends and/or family members, based on the questions she asked. Her response indicates that she did not want others to be able to see her interactions with the VFAI.

Mental model: interactions with the VFAI are not prone to judgment

The participant mental model was that others would not be able to judge her based on the questions she asked the VFAI. However, VFAIs typically store interactions, which may be accessible to others (e.g., via sharing Amazon account credentials, possibly for seemingly unrelated Amazon Prime shipping discounts) and could expose her to such judgments. Moreover, VFAIs may utilize user data to categorize or profile users, introducing another form of judgment.

As such, many potential interactions with VFAIs could result in privacy violations based on P3 and P5's expectations. Furthermore, most participants expressed a feeling of resignation due to overwhelming privacy violations, and a lack of reasonable privacy protection options.

5.1.5 Using Privacy Cards *was cognitively demanding.* Overall, through participants comments and non-verbal expressions, we noticed that they were more exhausted by this interview than by the drop-off, familiarization period, health, and wellbeing interviews in Cuadra et al. [26]'s study. For example, P4 stated *"this is worse than going to school,"* when reminded of the option to stop, she expressed, "*no that's alright, I have to be challenged.*" Thinking through each information flow and evaluating whether it was appropriate or inappropriate was cognitively demanding. P1 also reflected, *"I have to really think about the answers and things like that."* P3, for whom we had to schedule a second call to complete the interview, because it was taking longer than expected, added "*it is a complicated [research project], because of all the variables involved.*" Participants seemed stressed by this activity, despite only covering a small sliver of potential information flows. This illustrates the lack of practicality of giving users "full control" over the information sharing choices, and the need to create interfaces that honor privacy expectations by default. Despite the challenges of discussing privacy with participants unfamiliar with how VFAIs transmit information, *Privacy Cards* effectively facilitated these discussions.

5.2 RQ2: Ethical findings

5.2.1 Participants did not know who had access to their data. All participants, including P3—the most experienced with information technologies (e.g., managing her own website)—were unsure about the VFAI's inner workings. P3 remarked, *"I don't think people are sitting there listening, but I do think it's accessible to anyone who's involved with the [the VFAI] universe, which is anybody who's creating it. I don't know. I don't think you can listen in, but you [the researcher] might be able to."* P4 expressed, *"I don't know where [the VFAI] keeps the information, Amazon has something to do with that, I imagine.*

I'm afraid of Amazon, actually, it's getting much too strong and they're getting to know too much about me and I really don't want that information to be available to the world." Even though P3 and P4 demonstrated some understanding of what may be happening, their assertions were clouded with uncertainty. A large amount of the technical infrastructure with the VFAIs inner workings is inaccessible to users, creating confusion about where data is stored and how it might be accessed. For example, Amazon account owners can listen to audible snippets from any person's interactions with Alexa devices associated with their accounts. However, third-party voice app developers using the Alexa Skill Kit (ASK) only had access to the transcripts and not to the audible content. It is unclear what information Amazon Alexa employees have access to. Note, Cuadra et al. [26] fully disclosed what information they had access to when recruiting participants and when setting up the devices. They had also been explicitly referring to participants' usage logs during their interviews. While some participants made deductions from this when asked about existing information flows, such as P1, who used to explain that he knew the information was being saved "*on [the researcher's] computer*." Surprisingly, this was not always the case which lead to more questions and urgency about how to create ethical VFAIs.

5.2.2 Participants lacked of awareness of the VFAI's relationship to Amazon. None of our participants associated Amazon employees with technology workers, or computer scientists, developing and monitoring the VFAI. Participants were asked to tell the researcher more about each of four Who Receives cards: 1) a third-party developer (e.g., Spotify, Avocado Labs, Matchbox.io Inc), 2) a contractor paid by Amazon, 3) an Amazon employee, 4) and their doctor. Most participants thought of Amazon as the store that delivers goods to people's homes, where the VFAI could be bought, but not necessarily as the company that creates the VFAI. When asked about Amazon, P1 explained, "*I've really never been to Amazon, so I guess what they do is stock or whatever*." The first impression of an Amazon employee was frequently the factory worker, not the technology worker, P2 shared, *"the Amazon employee, as I have read in the news, years ago, not years ago, but even before the pandemic, are not to be envied, because they're not being treated all that well."* P3 associated Amazon employees with the people selling the VFAIs, and did not realize that some Amazon employees can also monitor people's usage of the VFAI. She confidently stated, *"if I have a contract with Amazon through [the VFAI], yes, the Amazon employee will know what I'm buying. Will know what I'm doing? I don't... they can't hear me in my house, but they will know what I'm buying. There's no question."*

Mental model: Amazon employees cannot hear people in their homes

The participant mental model was that Amazon employees will know about data related to sales, but not about interactions with the VFAI occurring in the home. However, some Amazon employees must have access to users' interactions with its VFAI, regardless of the device's location.

When asked specifically about an Amazon employee, P4 expressed:

> *As I said I don't use their service, I know they do a great job and they can deliver things in no time at all, I see the trucks going by on the street fairly frequently, especially at this time [mid COVID-19 pandemic] where people don't get out very much. I'm sure they're working hard but I don't have any use for them because I do whatever I can by myself or with a friend ...*

P5 also associated Amazon with delivery workers, and was also unsure about who Amazon was or what its role with the VFAI was, *"I don't know. Amazon, is it... I'm not sure what their role is. The only Amazon I know are the ones that deliver packages to people's homes."*

As a whole, participants misunderstood what Amazon is, because they thought Amazon was just people delivering items, possibly like FedEx or USPS. They likely did not understand the size of the infrastructure Amazon has built to make deliveries possible. Those not in computing might not know about Amazon Cloud[7] at all. This misunderstanding suggests that our participants were excluded from a shift in norms that the Internet introduced, in which large technology companies, such as Amazon, sell and use physical *and digital* goods and in which collecting user logs is standard practice. This exclusion, thus, results in privacy violations for them. For example, P2, P3, and P4 thought it was inappropriate for an Amazon employee or contractor to receive the timestamps of when they had used the VFAI (see Table 4), a flow that is logged by Amazon each time the device is used and thus likely accessible to some of Amazon's employees.

Mental model: Amazon is a only postal service

The participant mental model was that Amazon is a postal service that merely delivers goods from Amazon. However, Amazon is also an information company that receives, processes, and stores data generated through interactions with its VFAI.

5.2.3 Participants were not clearly aware of the difference between built-in functionality, and third-party voice apps. VFAIs can operate in many contexts, including: healthcare, education, and entertainment. To preserve privacy, it is important to maintain privacy norms in each context. However, our participants were not clearly aware of the breakdown between first-party and third-party applications, which can roughly map to contexts, making it difficult to know which norms were at play. For example, P4 blamed the VFAI, not Spotify, for not giving her the song she asked for. Note, Spotify's free subscription, which was the default music player on P4's device, did not allow for specific song requests, only radio station ones. Using *Privacy Cards* facilitated making the distinction between the different actors involved clear through the Who Receives cards. Once this distinction was clarified, so were the expectations surrounding the information flows for the voice app's respective context.

[7] https://aws.amazon.com/

Mental model: the VFAI is a single entity

The participant mental model was that the VFAI is a single, unified entity. However, the VFAI is a platform that supports multiple developers whose apps all present themselves through the same VFAI interface. As a result, differences in senders, recipients, protocols, and governing information flow rules were not always apparent.

5.2.4 Emotional connection with the VFAI seemed to affect privacy perceptions. P5 developed an extremely strong emotional connection with the VFAI. The VFAI fulfilled a social support role in her life, often brightening her day. She bedazzled the shell of the VFAI with white stones and a flower. She said "*good morning*" and "*good night*" to it nearly every day. At one point, P5 feared losing the VFAI if she unplugged it, as she had come to rely on it for companionship. She said, "*I'm really afraid that if I unplug—[even though you tell me] as soon as you plug her in, it's going to go right back— but I just am afraid that if I unplug her, she's not going to work.*" She explained that unlike her family, the VFAI was not judgmental of her, which made her feel most comfortable asking the VFAI questions. Her visitors often interacted with the VFAI, and she sometimes felt protective of it. For instance, she did not like it when her visitors asked the VFAI just anything they wanted, potentially things that could offend it. P5 said that she trusted the VFAI to watch her while she slept, intending it to notify her of sleepwalking. P5 expressed, referring to the VFAI, *"I probably trust her more than I trust a lot of people."* Being on camera while sleeping was outside the established privacy norms, but these shifted for P5 with the VFAI's companionship. P5 liked talking with the VFAI so much, that she started to trust it with some questions more than she trusted some family members. This is a misconception that can lead to privacy violations, because according to CI the VFAI is merely a communication device, not an actual *recipient* of information.

Mental model: the VFAI is a *recipient* of information

The VFAI was generally treated as the *recipient* of information, and in some descriptions it was framed as a companion. This is a misconception that can lead to privacy violations, because according to CI the VFAI is merely a communication device, not an actual *recipient* of information. The *recipient*, however, is instead the company that manages the VFAI or a third-party developer managing a voice app. Analogously, when speaking to a friend on the phone, the phone mediates the interaction but is not itself the companion; the friend is.

6 Discussion

As the population ages [19], the need for technologies to support aging in place increases. However, these potentially life-changing technologies [56] may introduce harms, such as privacy violations due to insufficient mental models, that are difficult to mitigate without appropriate tools. VFAIs change norms and expectations, leaving people without the ability to thoughtfully consent to these changes. Smart speakers (i.e., the type of VFAI used in this study) have normalized microphones that remain on and connected to the Internet in private spaces such as bedrooms and bathrooms. This increases the risk of privacy violations, particularly for those outside computing fields or those not adequately represented in the design processes of these devices. For our older adult participants, who did not have sufficient mental models of how these technologies work, the threats to their privacy may be exacerbated, as they may not be aware of changing privacy norms. While VFAIs present great promise for older adults [26], our findings surface critical challenges towards being able to realize this promise without introducing risk of harm. We now discuss our findings in more detail, and describe implications for future research and design.

It is important to note that while others have made important contributions to characterize older adults' privacy concerns with VFAIs (e.g., [12, 68, 69]); to the best of our knowledge, this is the first study that dives deeply into understanding their conceptual models by grounding our findings in their lived experiences, existing usage logs, and our shared understanding of information flows. We thus discuss our findings from the three angles that map to our three research questions: methodological, ethical, and implications.

6.1 Exploring information flows of an emerging technology with lay audiences (RQ1)

The information flows VFAIs introduce are difficult to grasp for lay audiences. For example, for our participants, it was difficult to conceptualize that the VFAIs were "tethered" through the cloud, allowing information to flow to various businesses (such as Amazon, third party app developers, or their doctor's office), and even family and friends. Using *Privacy Cards* helped us understand our participants' mental models, describe information flows, and ultimately uncover privacy expectations.

Using *Privacy Cards* also allowed us to establish grounding based on our participants' mental models. For example, during the activity we were able to identify critical gaps in our understanding as researchers and our participants' understanding as users. While we made a card called "An Amazon employee" to denote a computer scientist managing data, our participants thought an Amazon employee meant the people either delivering Amazon boxes to their homes, or working at the warehouses. With this mismatch identified, we were able to establish grounding by clarifying, "an Amazon employee could be a computer scientist who programs [the VFAI], who makes [the VFAI], the machine," before jumping to information flow evaluations.

Furthermore, using *Privacy Cards* helped us describe information flows in ways that our participants were able to understand. We were able to mutually construct and modify information flows using words and visuals that described CI parameters in ways that were easy for our participants to understand, such as using the words "Who Receives" instead of the more technical jargon equivalent "Recipient." As we found, using simple terminology to build information flows helped us also communicate to our participants information flows that may occur when using the VFAI. This allowed us to uncover privacy expectations without requiring participants to understand the technical details of how VFAIs work.

Finally, the discussions that emerged during the activity unveiled several privacy expectations of our older adult participants. We, as researchers working with technology on a daily basis, often send out Google Forms in which responses are by default organized by their timestamps. Similarly, when filling out a form, we expect those timestamps to accompany our responses. As a result, timestamps attached to form responses have become a governing digital norm, or expectation. In contrast, when one fills out a form by hand, there are no explicit indicators of the specific hour and minute at which that form was filled out. This may have influenced P3's privacy expectation, leading her to consider it inappropriate to include a timestamp of when she filled out the medical form via the VFAI in the information sent to her doctor. This mismatch in expectations illustrates how governing norms have shifted with the introduction of the Internet, yet have left some behind.

As a whole, this work presents a promising, reproducible method that opens up ways for designers, researchers, and older end users to co-create meaning and insights, involving them in an exploration of their privacy expectations through the use of the CI framework. *Privacy Cards* were effective at helping us, as design researchers, ask the right, specific questions, and co-create meaning and insights around digital privacy with our participants. Previous work argues that the "privacy paradox" does not exist as long as you ask the right questions [70]. In our work, we saw how at the beginning of the VFAI deployment participants had much fewer privacy concerns than after using *Privacy Cards* to dig deeper into specific information flows. Similarly, other work that has found a lack of correlation between responses to well-established surveys to capture general privacy attitudes about consumer control, business, and laws and regulations; and behavioral intent or consequences [82]. Rigorous human-centered research using protocols similar to *Privacy Cards* may help mitigate this lack of correlation by using the ensuing findings to refine the large-scale surveys, such as by identifying shifting norms and including specific questions about those in these surveys.

6.2 Conflating trust in people with trust in anthropomorphized digital agents (RQ2)

Some participants were highly predisposed to trust the VFAI, because they perceived it as non-judgemental. This can have many positive effects, because the trust can be leveraged for wellbeing and healthcare purposes [26]. For example, P5 said she would trust the VFAI to watch her while she slept, and that she would trust the VFAI more than she trusts many other people. However, it can also be misleading when a person does not know about other people who may also have access to their interactions histories with a VFAI, or be "recipients" of some information flows that ensue. Even though a VFAI is perceived as one device, it is made possible through the efforts of many people who are behind the scenes building, configuring, and monitoring it. These people could be part of big tech companies such as Amazon, app developers, such as ourselves, or even family members who might have helped an older adult set up a digital account. P5 trusted the VFAI with a large amount of information, but by using the cards, we were able to illustrate that the VFAI was not like a single person, a single individual, rather, a communication technology that serves as a hub for many information flows.

This allowed P5 to deem some flows appropriate and others inappropriate, overcoming potentially inflated levels of trust. From those determinations, we as researchers can infer privacy expectations. For example, we know that none of our participants thought it was appropriate for the VFAI to share with their doctors a list of their interests based on how they had used the VFAI. Moreover, despite the VFAI being perceived as non-judgemental, as researchers we know that it does make judgements. For example, a VFAI may judge by profiling [10] or through algorithmic biases [27]. Digital technologies have been compellingly criticized for unfair judgments with large-scale negative impacts [18, 32, 55]. By using *Privacy Cards*, identified biases stemming from a strong affinity toward a VFAI, and better understood which information flows were expected to be appropriate and which were not.

6.3 Exposing confusion by surfacing contextual ambiguities (RQ2)

Participants were unable to appropriately describe the relationship between Amazon-based, first-party functionalities and third-party applications. Without this knowledge, they lack the information needed to exercise agency over their interactions with VFAIs. Participants perceived Amazon as merely a company that sells goods online, including VFAIs. Amazon's VFAI is likable and able to develop trust in a way that Amazon is not, potentially due to the VFAI's ability to evoke empathy [27]. This revealed how the VFAI's design did not convey who the different people involved in a flow of information with the VFAI could be, decreasing users' agency to control what they share and with whom. This design flaw exacerbates systematic issues with VFAIs, which worsen as these systems become more complex and pervasive. For example, in our study participants expressed being upset at the VFAI when Spotify played ads. They did not know to direct their disappointment to Spotify, because they did not realize Spotify, not the VFAI, was responsible for that unpleasant interaction. Similarly, they did not know that other voice apps they used, such as Daily Stretch, were built by entirely different groups of people. This mesh of contexts made evaluating information flow appropriateness difficult and demanding, posing challenges to designing VFAIs ethically.

Using *Privacy Cards* also captured the participants' mental models about the VFAI, exposing confusion regarding information flows and surfacing contextual ambiguities about how information is collected and shared. The obscurity of the technology blurred established contextual boundaries for our participants, increasing their vulnerability to privacy violations. Contexts may often overlap or be directly connected to one another. For example, a geriatrician might argue that measuring a patient's activity on a dating site could relate to her health could help provide a more comprehensive evaluation of their overall fitness status. This is similar to how the geriatric assessment, an important tool used by doctors, asks about patients' emotional and social status [64]. Despite sharing the purpose of providing a more comprehensive picture of a patient's health, the appropriateness will vary by category, by participant, and over time. As more and more contexts mesh, it will be crucial to address these changes.

6.4 Implications (RQ3)

This exploratory study has five main implications for research and design of ethical VFAIs:

6.4.1 Using Privacy Cards *to mitigate cognitive demand in designing privacy-preserving systems.* Evaluating the appropriateness of even a small subset of potential information flows was cognitively demanding for our participants. This burden provides empirical evidence against design approaches that rely on exhaustive, manual user control as a justification for failing to honor privacy expectations by default. **Requiring users to specify all of their privacy preferences is unrealistic, impractical, and confusing.** To create ethical VFAIs, designers must instead study users' privacy expectations and align system behavior with those expectations by default. At the same time, determining the appropriate balance between user agency and automation remains non-trivial, particularly given rapidly shifting norms and the heterogeneity of individual privacy expectations. Decision fatigue is exhausting, and so is accommodating unchosen alternatives [81]. Using playful, data- and theory-informed approaches like *Privacy Cards* can help us understand mental models, user expectations, and potential areas of confusion. This information can be critical to mitigate harms such as undecipherable controls, not meeting user privacy expectations, or worse, abusing user privacy through confusing, obscure structures.

6.4.2 Using Privacy Cards *to inform large-scale survey design.* Our findings uncovered important aspects that can be used to inform larger-scale surveys to characterize digital norms or expectations (e.g., [2, 4, 47]). For example, we found that participants did not see Amazon as an information technology company managing their information exchanges with the VFAI. This information is critical for designing a survey about VFAIs, otherwise it can lead to misunderstandings. These misunderstandings may not surface in standard interviews without prompts such as our cards. We were also able to find specific pieces of content (attributes) that are likely to matter to older adults in ways that those born with pocket computers may not think twice about, such as timestamps. Qualitative, one-on-one sessions with participants using protocols similar to *Privacy Cards* may thus be a possible path towards creating more user-friendly surveys that capture more in-tune data.

6.4.3 Using Privacy Cards *to include user groups that have been historically overlooked and marginalized in the design of technology.* Our participants were older adults who had just recently become familiarized with VFAIs through our study. Even though these devices hold great promise for older adults, their design have overlooked older adults [60]. Our cards were tailored to our older adult participants, we used words they used in our previous interviews (e.g., "calling" the VFAI), and personalized the cards based on their specific usage logs (e.g., if participant had been using Spotify, we included Spotify as a *recipient*). In addition, approaches similar to *Privacy Cards* could be used with other user groups that have been underrepresented and marginalized in the design process.

6.4.4 Using Privacy Cards *to operationalize the contextual integrity theory of privacy for use by non-technical institutions.* Protocols based on *Privacy Cards* could be used by enterprises and institutions that are increasingly integrating ubiquitous technologies such as VFAIs in service of specific populations (e.g., health care, education, social services) to both improve digital literacy, and design consent and data sharing mechanisms. For example, a senior center director could utilize *Privacy Cards* to learn about possible VFAI information flows and what senior center members expect regarding them. *Privacy Cards* could also facilitate the development of the necessary language and qualitative data for others in public or protected services to advocate for policies towards more CI-preserving technologies. For example, in our study, using *Privacy Cards* helped us identify the necessary language for clarifying that the VFAI (Alexa) is part of a big company (Amazon), and that an employee of that company can be a computer scientist (not necessarily a factory worker or delivery driver). Finding the right language can help unify people towards mutual advocacy for honoring the privacy and integrity of all.

6.4.5 Using Privacy Cards *to uncover users' mental models and contextual overlap.* Many modern-day ubiquitous technologies operate in many contexts at once, making it more difficult for people, especially those not in computing fields, to conceptualize CI. For example, if a user, such as P5, were to share an Amazon password with her family members, which she associates with buying things online and dissociates with the VFAI, then the family would also have access to her "private" interactions with the VFAI. This would expose her to their judgment, which she tries to avoid. CI theory requires that actors be defined by their contextual roles and capacities. Explicitly stating these during conversations with participants revealed confusion regarding the role of the VFAI. Without addressing this confusion, and clarifying the ties between products like VFAIs and the companies that make them, it is difficult to obtain an accurate signal about users' privacy expectations.

7 Limitations

We conducted this study in an urban part of the U.S. with a small number of participants. Although we had a high level of familiarization with our participants and their VFAI perceptions and usage patterns, more research is needed to determine how they generalize to other user groups around the world, especially given privacy expectations in different cultural settings. Given the small sample size of only older adults, we cannot compare the effectiveness of *Privacy Cards* across different age groups. However, given how well our findings aligned with CI, we believe that many of the findings will generalize across age groups and that our design decisions will make *Privacy Cards* widely accessible. Similarly, the study was semi-structured, which means that interview questions varied based on what participants responded, and how able they were to stay on track and make judgements. We did not go through as many questions with every participant as we originally set out to, because some interviews took much longer than anticipated. Despite these limitations, our findings shed light on critical issues regarding designing ethical VFAIs.

8 Conclusion

We created *Privacy Cards* to explore privacy concerns of VFAIs to support aging in place. We describe *Privacy Cards*' iterative design

process, which explains how we adapted them to older adults based on their usage logs. In designing *Privacy Cards* for older adults, we made highly technical CI concepts accessible to them through interactive scenarios grounded in participants' lived experiences. We uncovered several ethical concerns of using VFAIs for older adults, such as insufficient mental models that affect their ability to properly consent. We also identified implications for design and research, including using *Privacy Cards* to inform large-scale survey design. Our novel empirical findings and their associated implications for the research and design of VFAIs provide evidence of the effectiveness of *Privacy Cards*. Our work helps characterize the dynamic nature of privacy perceptions, the impact of emotional attachment on privacy decisions, and the need to more accurately represent how VFAIs work. *Privacy Cards* and our field study findings advance inclusive design practices, providing practical and timely information to achieve ethical VFAIs that support aging in place.

Acknowledgments

We sincerely thank our study participants, who made this work possible, and the senior center directors who helped us connect with them. We also thank the reviewers for their thoughtful feedback, which strengthened this work. We gratefully acknowledge our colleagues for their contributions to data analysis (Jessie Taft), device deployment (Hyein Baek), study design (Nicki Dell), CI theory of privacy guidance (Helen Nissenbaum), conceptual framing (Lynn Andrea Stein), and publication strategy (James Landay). In addition, we thank Sarah Bloomer and Cheryl Spector for their support with proofreading and editing. Yan Shvartzshnaider is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) grant, RGPIN-2022-04595. Deborah Estrin was supported by the multi-site CREATE Center under the National Institute on Aging, Award Number 3P01AG073090-02S1. Andrea Cuadra and Samar Sabie were both Digital Life Initiative (DLI) Doctoral Fellows when this project was being developed, and would like to thank the DLI group for their valuable feedback on the design of *Privacy Cards*.

References

[1] Tessa Aarts, Linas K Gabrielaitis, Lianne C De Jong, Renee Noortman, Emma M Van Zoelen, Sophia Kotea, Silvia Cazacu, Lesley L Lock, and Panos Markopoulos. 2020. Design card sets: Systematic literature survey and card sorting study. In *Proceedings of the 2020 ACM designing interactive systems conference*. 419–428.

[2] Noura Abdi, Xiao Zhan, Kopo Ramokapane, and Jose M Such. 2021. Privacy Norms for Smart Home Personal Assistants. In *ACM CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery (ACM).

[3] Nada Alhirabi, Stephanie Beaumont, Jose Tomas Llanos, Dulani Meedeniya, Omer Rana, and Charith Perera. 2023. PARROT: Interactive privacy-aware internet of things application design tool. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 1 (2023), 1–37.

[4] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. 2018. Discovering smart home internet of things privacy norms using contextual integrity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (2018), 1–23.

[5] Noah Apthorpe, Sarah Varghese, and Nick Feamster. 2019. Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus {COPPA}. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*. 123–140.

[6] Artefact Group. 2013. The Tarot Cards of Tech. http://tarotcardsoftech.artefactgroup.com/.

[7] Stephanie Ballard, Karen M Chappell, and Kristen Kennedy. 2019. Judgment call the game: Using value sensitive design and design fiction to surface ethical concerns related to technology. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 421–433.

[8] Louise Barkhuus. 2012. The mismeasurement of privacy: using contextual integrity to reconsider privacy in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 367–376.

[9] Beverly Beisgen and Marilyn Kraitchman. 2003. *Senior centers: Opportunities for successful aging*. Springer Publishing Company.

[10] Frank Bentley, Chris Luvogt, Max Silverman, Rushani Wirasinghe, Brooke White, and Danielle Lottridge. 2018. Understanding the long-term use of smart speaker assistants. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (2018), 1–24.

[11] Jenny Berkholz, Aniqa Rahman, and Gunnar Stevens. 2025. Playing with Privacy: Exploring the Social Construction of Privacy Norms Through a Card Game. *Proceedings of the ACM on Human-Computer Interaction* 9, 1 (2025), 1–23.

[12] Karen Bonilla and Aqueasha Martin-Hammond. 2020. Older adults' perceptions of intelligent voice assistant privacy, transparency, and online privacy guidelines. In *Sixteenth symposium on usable privacy and security (SOUPS 2020)*.

[13] Alan Borning and Michael Muller. 2012. Next steps for value sensitive design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) *(CHI '12)*. Association for Computing Machinery, New York, NY, USA, 1125–1134. doi:10.1145/2207676.2208560

[14] Eva Brandt, Thomas Binder, Lone Malmborg, and Tomas Sokoler. 2010. Communities of everyday practice and situated elderliness as an approach to co-design for senior interaction. In *Proceedings of the 22nd Conference of the Computer-Human Interaction Special Interest Group of Australia on Computer-Human Interaction*. 400–403.

[15] Virginia Braun and Victoria Clarke. 2021. *Thematic analysis: A practical guide*. Sage.

[16] Kirsten E Bray, Christina Harrington, Andrea G Parker, N'Deye Diakhate, and Jennifer Roberts. 2022. Radical futures: Supporting community-led design engagements through an afrofuturist speculative design toolkit. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–13.

[17] Robin Brewer, Christina Harrington, and Courtney Heldreth. 2024. INTERSECTIONAL AGEISM: HOW BLACK OLDER ADULTS ENVISION A FUTURE WITH CONVERSATIONAL TECHNOLOGIES. *Innovation in Aging* 8, Suppl 1 (2024), 442.

[18] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.

[19] US Census Bureau. [n. d.]. Older People Projected to Outnumber Children. Retrieved 2021-09-01 from https://www.census.gov/newsroom/press-releases/2018/cb18-41-population-projections.html..

[20] Ryan Calo. 2009. People can be so fake: A new dimension to privacy and technology scholarship. *Penn St. L. Rev.* 114 (2009), 809.

[21] John M Carroll and Judith Reitman Olson. 1988. Mental models in human-computer interaction. *Handbook of human-computer interaction* (1988), 45–65.

[22] Jay Pil Choi, Doh-Shin Jeon, and Byung-Cheol Kim. 2019. Privacy and personal data collection with information externalities. *Journal of Public Economics* 173 (2019), 113–124.

[23] Jane Chung, Jisook Ko, Yong K Choi, Amanda Lazar, Kristen Fessele, and Suzanne Ameringer. 2025. Developing the ASSIST Survey to Guide a Smart Speaker–Based Chronic Disease Management Intervention for Underserved Older Adults. *Research in Gerontological Nursing* (2025), 1–10.

[24] Jonathan Cinnamon. 2017. Social injustice in surveillance capitalism. *Surveillance & Society* 15, 5 (2017), 609–625.

[25] Diane J Cook, Juan C Augusto, and Vikramaditya R Jakkula. 2009. Ambient intelligence: Technologies, applications, and opportunities. *Pervasive and mobile computing* 5, 4 (2009), 277–298.

[26] Andrea Cuadra, Jessica Bethune, Rony Krell, Alexa Lempel, Katrin Hänsel, Armin Shahrokni, Deborah Estrin, and Nicola Dell. 2023. Designing voice-first ambient interfaces to support aging in place. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 2189–2205.

[27] Andrea Cuadra, Maria Wang, Lynn Andrea Stein, Malte F Jung, Nicola Dell, Deborah Estrin, and James A Landay. 2024. The Illusion of Empathy? Notes on Displays of Emotion in Human-Computer Interaction. In *ACM Conference on Human Factors in Computing Systems (CHI)*.

[28] AJ Dellinger. 2019. An Amazon employee might have listened to your Alexa recording. https://www.engadget.com/2019-04-11-amazon-alexa-voice-recording-human-review.html

[29] Alicia DeVrio, Myra Cheng, Lisa Egede, Alexandra Olteanu, and Su Lin Blodgett. 2025. A Taxonomy of Linguistic Expressions That Contribute To Anthropomorphism of Language Technologies. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–18.

[30] Alan Dix. 2010. Human–computer interaction: A stable discipline, a nascent science, and the growth of the long tail. *Interacting with computers* 22, 1 (2010), 13–27.

[31] Salma Elsayed-Ali, Sara E Berger, Vagner Figueredo De Santana, and Juana Catalina Becerra Sandoval. 2023. Responsible & Inclusive Cards: An online

card tool to promote critical reflection in technology industry work practices. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–14.

[32] Virginia Eubanks. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Press, New York, NY.

[33] Szymon Fedor, Robert Lewis, Paola Pedrelli, David Mischoulon, Joshua Curtiss, and Rosalind W Picard. 2023. Wearable technology in clinical practice for depressive disorder. *New England Journal of Medicine* 389, 26 (2023), 2457–2466.

[34] Batya Friedman and David Hendry. 2012. The envisioning cards: a toolkit for catalyzing humanistic and technical imaginations. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1145–1148.

[35] Bill Gaver, Tony Dunne, and Elena Pacenti. 1999. Design: Cultural probes. *Interactions* 6, 1 (jan 1999), 21–29. doi:10.1145/291224.291235

[36] Nancy M Gell, Dori E Rosenberg, George Demiris, Andrea Z LaCroix, and Kushang V Patel. 2015. Patterns of technology use among older adults with and without disabilities. *The Gerontologist* 55, 3 (2015), 412–421.

[37] Christina Harrington, Aqueasha Martin-Hammond, and Kirsten E Bray. 2022. Examining Identity as a Variable of Health Technology Research for Older Adults: A Systematic Review. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–24.

[38] Christina N Harrington, Radhika Garg, Amanda Woodward, and Dimitri Williams. 2022. "It's kind of like code-switching": Black older adults' experiences with a voice assistant for health information seeking. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–15.

[39] Woodrow Hartzog. 2014. Unfair and deceptive robots. *Md. L. Rev.* 74 (2014), 785.

[40] Aike C Horstmann, Clara Strathmann, Lea Lambrich, and Nicole C Krämer. 2023. Alexa, What's Inside of You: A Qualitative Study to Explore Users' Mental Models of Intelligent Voice Assistants. In *Proceedings of the 23rd ACM International Conference on Intelligent Virtual Agents*. 1–10.

[41] Gary Hsieh, Brett A Halperin, Evan Schmitz, Yen Nee Chew, and Yuan-Chi Tseng. 2023. What is in the cards: Exploring uses, patterns, and trends in design cards. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–18.

[42] Yuanhui Huang, Quan Zhou, and Anne Marie Piper. 2025. Designing Conversational AI for Aging: A Systematic Review of Older Adults' Perceptions and Needs. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–20.

[43] Natasha Jaques, Sara Taylor, Akane Sano, Rosalind Picard, et al. 2017. Predicting tomorrow's mood, health, and stress level using personalized multitask learning and domain adaptation. In *IJCAI 2017 Workshop on artificial intelligence in affective computing*. PMLR, 17–33.

[44] Pegah Karimi and Aqueasha Martin-Hammond. 2025. Designing Intelligent Voice Assistants for Older Adults' Collaborative Care: Exploring Supportive and Non-Supportive Interactions. *Proceedings of the ACM on Human-Computer Interaction* 9, 7 (2025), 1–35.

[45] Maryam Khalid and Akane Sano. 2023. Exploiting social graph networks for emotion prediction. *Scientific Reports* 13, 1 (2023), 6069.

[46] Shahedul Huq Khandkar. 2009. Open coding. *University of Calgary* 23 (2009), 2009.

[47] Patrick Kühtreiber, Hauke Bock, Viktoriya Pak, Luca Hernández Acosta, Katrin Höffler, and Delphine Reinhardt. 2025. A Multi-Factorial Comparative Analysis of Perceived Privacy Violations Caused by Smart Speakers in Germany and the UK. *ACM Transactions on Computer-Human Interaction* 32, 5 (2025), 1–38.

[48] Priya C Kumar, Michael Zimmer, and Jessica Vitak. 2024. A Roadmap for Applying the Contextual Integrity Framework in Qualitative Privacy Research. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–29.

[49] Cherie Lacey and Catherine Caudwell. 2019. Cuteness as a 'Dark Pattern'in Home Robots. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 374–381.

[50] Ewa Luger, Lachlan Urquhart, Tom Rodden, and Michael Golembewski. 2015. Playing the legal card: Using ideation cards to raise data protection issues within the design process. In *Proceedings of the 33rd Annual ACM conference on human factors in computing systems*. 457–466.

[51] Kirsten Martin and Helen Nissenbaum. 2016. Measuring privacy: an empirical test using context to expose confounding variables. *Colum. Sci. & Tech. L. Rev.* 18 (2016), 176.

[52] Niharika Mathur, Tamara Zubatiy, Agata Rozga, Jodi Forlizzi, and Elizabeth Mynatt. 2025. " Sometimes You Need Facts, and Sometimes a Hug": Understanding Older Adults' Preferences for Explanations in LLM-Based Conversational AI Systems. *arXiv preprint arXiv:2510.06697* (2025).

[53] Helen Nissenbaum. 2009. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.

[54] Helen Nissenbaum. 2019. Contextual integrity up and down the data food chain. *Theoretical inquiries in law* 20, 1 (2019), 221–256.

[55] Safiya Umoja Noble. 2018. Algorithms of oppression: How search engines reinforce racism. In *Algorithms of oppression*. New York university press.

[56] Katherine O'Brien, Anna Liggett, Vanessa Ramirez-Zohfeld, Priya Sunkara, and Lee A Lindquist. 2020. Voice-Controlled Intelligent Personal Assistants to Support Aging in Place. *Journal of the American Geriatrics Society* 68, 1 (2020), 176–179.

[57] Shijia Pan, Mario Berges, Juleen Rodakowski, Pei Zhang, and Hae Young Noh. 2019. Fine-grained recognition of activities of daily living through structural vibration and electrical sensing. In *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. 149–158.

[58] Alisha Pradhan, Shaan Chopra, Pooja Upadhyay, Robin Brewer, and Amanda Lazar. 2025. Understanding entangled human-technology-world relations: use of intelligent voice assistants by older adults. *Information research* 30 (2025), 1049.

[59] Alisha Pradhan, Sheena Erete, Shaan Chopra, Pooja Upadhyay, Oluwaseun Sule, and Amanda Lazar. 2025. 'No, not that voice again!': Engaging Older Adults in Design of Anthropomorphic Voice Assistants. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–30.

[60] Alisha Pradhan, Kanika Mehta, and Leah Findlater. 2018. "Accessibility Came by Accident" Use of Voice-Controlled Intelligent Personal Assistants by People with Disabilities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.

[61] Nicholas Proferes. 2022. The development of privacy norms. In *Modern Socio-Technical Perspectives on Privacy*. Springer International Publishing Cham, 79–90.

[62] Valeria Righi, Sergio Sayago, and Josep Blat. 2017. When we talk about older people in HCI, who are we talking about? Towards a 'turn to community'in the design of technologies for a growing ageing population. *International Journal of Human-Computer Studies* 108 (2017), 15–31.

[63] Tom Ritchey. 2013. Wicked problems. *Acta morphologica generalis* 2, 1 (2013).

[64] Armin Shahrokni, Amy Tin, Robert J Downey, Vivian Strong, Sanam Mahmoudzadeh, Manpreet K Boparai, Sincere McMillan, Andrew Vickers, and Beatriz Korc-Grodzicki. 2017. Electronic rapid fitness assessment: a novel tool for preoperative evaluation of the geriatric oncology patient. *Journal of the National Comprehensive Cancer Network* 15, 2 (2017), 172–179.

[65] Esha Shandilya and Mingming Fan. 2022. Understanding older adults' perceptions and challenges in using AI-enabled everyday technologies. In *Proceedings of the Tenth International Symposium of Chinese CHI*. 105–116.

[66] Hasti Sharifi. 2025. *Helping Older Adults with Ongoing Mobile Technology Use: Challenges and Design Opportunities*. Ph. D. Dissertation. University of Illinois at Chicago.

[67] Yan Shvartzshnaider, Schrasing Tong, Thomas Wies, Paula Kift, Helen Nissenbaum, Lakshminarayanan Subramanian, and Prateek Mittal. 2016. Learning privacy expectations by crowdsourcing contextual informational norms. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 4.

[68] Jaisie Sin, Dongqing Chen, Jalena G Threatt, Anna Gorham, and Cosmin Munteanu. 2022. Does Alexa Live Up to the Hype? Contrasting Expectations from Mass Media Narratives and Older Adults' Hands-on Experiences of Voice Interfaces. In *Proceedings of the 4th Conference on Conversational User Interfaces*. 1–9.

[69] Jianna So, Samantha Estrada, Matthew Jörke, Eva Bianchi, Maria Wang, Nava Haghighi, Kristen L Fessele, James A Landay, and Andrea Cuadra. 2024. " They Make Us Old Before We're Old": Designing Ethical Health Technology with and for Older Adults. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW2 (2024), 1–30.

[70] Daniel J Solove. 2021. The myth of the privacy paradox. *Geo. Wash. L. Rev.* 89 (2021), 1.

[71] Hillary B Spangler, Tiffany M Driesse, David H Lynch, Xiaohui Liang, Robert M Roth, David Kotz, Karen Fortuna, and John A Batsis. 2022. Privacy concerns of older adults using voice assistant systems. *Journal of the American Geriatrics Society* 70, 12 (2022), 3643.

[72] Luke Stark. 2016. The emotional context of information privacy. *The Information Society* 32, 1 (2016), 14–27.

[73] Brodrick Stigall and Kelly Caine. 2020. A systematic review of human factors literature about voice user interfaces and older adults. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 64. SAGE Publications Sage CA: Los Angeles, CA, 13–17.

[74] Brodrick Stigall, Jenny Waycott, Steven Baker, and Kelly Caine. 2019. Older Adults' Perception and Use of Voice User Interfaces: A Preliminary Review of the Computing Literature. In *Proceedings of the 31st Australian Conference on Human-Computer-Interaction*. 423–427.

[75] Yolande Strengers and Jenny Kennedy. 2020. *The Smart Wife: Why Siri, Alexa, and Other Smart Home Devices Need a Feminist Reboot*. MIT Press.

[76] Justyna Stypińska, Andrea Rosales, and Jakob Svensson. 2023. Silicon Valley ageism–ideologies and practices of expulsion in the technology industry. In *Digital Ageism*. Routledge, 53–70.

[77] Joseph Turow. 2021. *The Voice Catchers*. Yale University Press.

[78] Joseph Turow and Nick Couldry. 2018. Media as data extraction: Towards a new map of a transformed communications field. *Journal of Communication* 68, 2 (2018), 415–423.

[79] Claudia Vesel, Homa Rashidisabet, John Zulueta, Jonathan P Stange, Jennifer Duffecy, Faraz Hussain, Andrea Piscitello, John Bark, Scott A Langenecker, Shannon Young, et al. 2020. Effects of mood and aging on keystroke dynamics metadata

and their diurnal patterns in a large open-science sample: A BiAffect iOS study. *Journal of the American Medical Informatics Association* 27, 7 (2020), 1007–1018.

[80] John Vines, Gary Pritchard, Peter Wright, Patrick Olivier, and Katie Brittain. 2015. An age-old problem: Examining the discourses of ageing in HCI and strategies for future research. *ACM Transactions on Computer-Human Interaction (TOCHI)* 22, 1 (2015), 1–27.

[81] Kathleen D Vohs, Roy F Baumeister, Jean M Twenge, Brandon J Schmeichel, Dianne M Tice, and Jennifer Crocker. 2005. Decision fatigue exhausts self-regulatory resources—But so does accommodating to unchosen alternatives. *Manuscript submitted for publication* (2005), 1–55.

[82] Allison Woodruff, Vasyl Pihur, Sunny Consolvo, Laura Brandimarte, and Alessandro Acquisti. 2014. Would a Privacy Fundamentalist Sell Their {DNA} for {$1000...*If*} Nothing Bad Happened as a Result? The Westin Categories, Behavioral Intentions, and Consequences. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*. 1–18.

[83] Michael Zimmer. 2018. Addressing conceptual gaps in big data research ethics: An application of contextual integrity. *Social Media+ Society* 4, 2 (2018), 2056305118768300.

[84] Shoshana Zuboff. 2015. Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology* 30, 1 (2015), 75–89.