Should I Agree? Delegating Consent Decisions Beyond the Individual

Bettina Nissen

Design Informatics
University of Edinburgh, UK
bettina.nissen@ed.ac.uk

Victoria Neumann

Lancaster University, UK
v.neumann@lancaster.ac.uk

Mateusz Mikusz

Lancaster University, UK
m.mikusz@lancaster.ac.uk

Rory Gianni

Design Informatics
University of Edinburgh, UK
r.gianni@ed.ac.uk

Sarah Clinch

School of Computer Science,
University of Manchester, UK
sarah.clinch@manchester.ac.uk

Chris Speed

Design Informatics
University of Edinburgh, UK
c.speed@ed.ac.uk

Nigel Davies

Lancaster University, UK
n.a.davies@lancaster.ac.uk

ABSTRACT

Obtaining meaningful user consent is increasingly problematic in a world of numerous, heterogeneous digital services. Current approaches (e.g. agreeing to Terms and Conditions) are rooted in the idea of individual control despite growing evidence that users do not (or cannot) exercise such control in informed ways. We consider an alternative approach whereby users can opt to delegate consent decisions to an ecosystem of third-parties including friends, experts, groups and AI entities. We present the results of a study that used a technology probe at a large festival to explore initial public responses to this reframing – focusing on when and to whom users would delegate such decisions. The results reveal substantial public interest in delegating consent and identify differing preferences depending on the privacy context, highlighting the need for alternative decision mechanisms beyond the current focus on individual choice.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI); Field studies;**

KEYWORDS

Consent; Delegation; Privacy; Permission Management; Technology Probe; Design

ACM Reference format:

Bettina Nissen, Victoria Neumann, Mateusz Mikusz, Rory Gianni, Sarah Clinch, Chris Speed, and Nigel Davies. 2019. Should I Agree? Delegating Consent Decisions Beyond the Individual. In *Proceedings of CHI Conference on Human Factors in Computing Systems Proceedings, Glasgow, Scotland UK, May 4–9, 2019 (CHI 2019)*, 13 pages. <https://doi.org/10.1145/3290605.3300745>

1 INTRODUCTION

Technologies such as the Internet of Things and ubiquitous mobile devices increase the frequency, volume and sensitivity of data collected, leading to ever more complex systems in which users are expected to understand, navigate and control the flow of their data. Current mechanisms designed to obtain consent from individuals such as offering a binary choice of acceptance/rejection of multi-page Terms and Conditions (T&Cs), have been widely criticised with specific challenges identified by scholars including a lack of control regarding the granularity of data access [58], the overwhelming and impractical burden put on the individual to consent to every possible data access [47–49], and the pitfalls of approaches such as blanket consent [33]. While theoretical and methodological debates about privacy and consent are not new, advances in the reach and impact of technology as well as changes in policy (e.g. the EU's GDPR [64]) have renewed critique of existing privacy models, leading some to claim we live in non-consensual worlds [6, 9, 24].

In this paper, we explore an alternative approach to consent in which users can choose to delegate consent decisions to an ecosystem of third-parties including friends, experts,

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

© 2019 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-5970-2/19/05.
<https://doi.org/10.1145/3290605.3300745>

groups and AI entities. In particular, we aim to develop an initial understanding of *when* and *to whom* users would delegate consent decisions. We present the results of a design research study using a technology probe [37] at a large international festival to explore initial public responses to this reframing of consent as delegation. Quantitative data describing 565 users' interactions is combined with qualitative data from 44 participant questionnaires and researcher observations. Our findings highlight the shortcomings of existing consent models, provide evidence of public acceptance of the notion of delegating consent and suggest the need for providing users with flexibility in delegation options for consent decisions.

Figure 1: Visitor interacting with *Trustball*

2 UNDERSTANDING PRIVACY AND CONSENT

Consent is an important field of theoretical and empirical inquiry in many disciplines. In relation to information technologies, it is often associated with privacy and data sharing and this paper focuses on consent as it relates to these issues.

Theoretical Framing

Themes of privacy, control and autonomy are rooted in libertarian thought schools from the enlightenment [6, 51]. At its core, there is an isolated individual self (i.e. free from outside influences) that is able to choose and make autonomous decisions regarding what to disclose and what not. The subsequent decisions are mainly conceptualised in terms of rational choice, which is noticeable in the popular framing of cost-benefit analysis [77]. Consent in this sense is the guard, which protects an inner, private realm from an outside, public realm. However, this view has been widely criticised [69]. Moving away from the Bourgeois practices of privacy as retreat of the individual from the public or social life, 20th century privacy discourses moved to practices of information control of multiple instances of the self in a relational performance [52, 62]. However, this is also not sufficient for the needs of a networked and entangled self of the 21st century. What's at stake is not the revealing of an "inner truth"

about an individual but rather connectivity to others in the digital, connected society [52]. As Bechmann notes, "individual privacy is downplayed as a result of the click-wrap agreement culture on the internet." [6]. This brings attention to consent and the limitations and costs of user-centred control approaches [5, 53], where privacy is conceptualised as part of self-management practices. This is often highlighted as the paradox between users' attitudes towards privacy and their disclosure behaviour in practice [75]. Drawing on Nissenbaum's analysis of the privacy paradox [56, 57], Bechman [6] argues that privacy exists in a contradictory relationship between the users' desire for data protection privacy versus the necessity for services and their convenience. This often results in what has been discussed as "preference uncertainty" [34] influenced by many factors which may go beyond an individual's understanding or perception of control. Nissenbaum [56, 57] and Solove [71] stress the importance of the "consent dilemma" when discussing informed consent and emphasise that privacy and consent decisions have to be seen as contextual. When shifting the focus away from data permission management, consent is also an issue in other areas. In legal and medical fields, consent is the basis of making contracts and exercising freedoms and rights, such as autonomy, while balancing it with questions of liability [25]. However, delegation of consent (e.g. to legal guardians, family members or business partners) is more common than in digital domains [23, 35, 38]. Further empirical work for delegating other decisions in related areas are, e.g., supporting people with visual impairments [10] or behaviour change plans [2].

Consent, Trust and Intermediaries

Privacy intermediaries and recommendation-based systems for consent have been explored in [1, 12, 81] including for collecting consent decisions under a single interface [41, 49], simplifying privacy policies [7] or automating decision making [4, 41]. From our review, we identified four main categories of approach: First, *community or crowd-based systems* are either community-based projects such as *Terms of Service; Didn't Read* (ToS;DR) [63] or privacy intermediaries that use crowd-sourcing approaches to support users to make data permission decisions [42, 45]. Second, *machine learning systems* and automated agents are a popular approach to privacy intermediaries, some focusing on simplifying privacy policies or T&Cs to fix the "biggest lie on the internet" [11, 48, 59] while other works use ontology or agent-based approaches [4, 72] to support users maintain their own [43] and others' privacy [27]. Third, *expert-based systems* are approaches, services or organisations that assist decision-making by providing users with expert recommendations on privacy settings [3, 65, 66]. Finally, other systems use *nominated individuals* to review and authenticate

online access, such as tools that obtain verifiable parental consent [15]. In many of these studies, trust was recognised as an important element of the acceptance of recommendations. Trust influences how people perceive information [45], evaluate risk on information disclosure [8, 39], or decide on attitudes towards secondary data use [40]. The element of trust seems to reduce the need for control. As Libaque-Saenz et al. [44] put it: “[...] in these trust-based relationships, trustors are motivated to rely on proxy-control rather than self-control.”

Reframing Consent As Delegation

The above review shows how the discourse on privacy intermediaries follows an individualised, user-centric approach, which falls into the realm of informational self-protection [60]. It is based on the framing that individuals have the responsibility to protect their privacy in the face of massive power and informational asymmetries between users and service providers [61]. The lack of resources (temporal, social, cognitive, or material) of single individuals to cope with these imbalances are amplified by the demand to manage complex issues through one-click consent. The result is that research has repeatedly shown that users do not actually read T&Cs and hence any notion of having obtained “informed consent” is problematic [55]. In this paper we explore the idea of reframing consent as delegation practice – transforming the question for users from “do I agree with a set of T&Cs” into “who can tell me whether I should agree to these T&Cs”.

3 STUDY DESIGN AND METHOD

Approach

To gain insights into public opinion and behaviour regarding consent and the willingness to consider delegation, we took a design research approach. While digital surveys are traditionally used, they can often be limited or biased [21] and lack potential for socially rich engagement with a subject matter. Our approach was intended to allow group interactions (reflecting on privacy socialisation and group norms), to offer opportunities for observation and more mutual exchange than other data collection approaches. A series of recent design research studies in HCI have adopted novel approaches to gather contextual information from audiences. For example, interactive versions of feedback questionnaires have been explored for a variety of audiences and purposes including to support care organisations [22], activist movements [79] and community engagement [74]. These interactive voting, surveying or feedback devices range in nature from mostly digital [22] to hybrid [79] or more physical [74] forms. Increasingly these interactive surveys or feedback devices have been adopted in HCI research in more tangible forms of questionnaires [28, 29, 31, 32] which have been used

to collect data from public audiences drawing on the material properties and affordances of physical artefacts to engage different audiences in the topic of the study, e.g. surveying fablab visitors [32], children [28] or elderly audiences [67]. In public settings, tangible questionnaires offer the possibility to be physically situated within their context and to call on people’s curiosity and desire to investigate a unique or novel artefact [36]. In our case, we intentionally aimed to remove the questionnaire from the mundane, textual context of consent and user agreements, and introduce a more experiential and reflective method to motivate public audiences to participate. In this sense, our tangible questionnaire takes on characteristics of a Technology Probe [37], combining social science goals of collecting data with HCI and design goals of exploring new technologies in real-life settings intended to inspire novel ideas and ways of thinking. The term probes in HCI has been widely adopted from its original cultural probes [30] through to design probes [80] and is increasingly used to describe a diverse variety of methods and practices [13] in co-design or participatory design settings. Technology probes are considered extensions of these forms of probes because their underlying technology simultaneously collects data while providing insights into current practices or systems [37]. We drew on these methods to invite a large public audience who may not usually participate in research studies to engage with our research topic in a playful manner while also challenging them to reflect on and rethink their consent practices.

Following these methods, the aim of our probe and this research study was not solely to gather data about user’s attitudes towards privacy and consent but to create a condensed consent experience that exemplifies signing up to a new app or service beyond acting only as questionnaire. An online or emailed questionnaire would remove questions from the contextual nature of consent and its experience, incentive or motivation of making user agreements. We aimed to embody the “experimental and subversive nature of the original probes” [13, p.1081] in our design to create an experience that may recreate a desire usually present when confronted with T&Cs. Many people agree to T&Cs when signing up to a new app or service and are engaged in a process of use rather than reflectively considering the implications of consent. In addition, these situations can often come with social pressure, e.g. a friend wants you to use a new app or you would like to read a news article so you quickly accept T&Cs in that moment. We aimed to incorporate this contextual nature of being ‘put on the spot’ in our experiential survey. We hoped that this colourful and experiential design engaged audiences through a similar desire to play (and agree) with the situational, ad hoc user choices that are common in the contextual reality of user agreement interactions.

Designing Trustball

Following this methodology, we developed an interactive, tangible device which would manifest our questions for the audience in an exciting and curious manner. Because the context for the installation was an international cultural event lasting several weeks, we decided against a low fidelity design which is often common in probes. To engage and draw the attention of public footfall, we opted to develop a high-quality designed artefact entitled *Trustball* that we hoped would encourage visitors to ‘play’ our questionnaire.

Figure 2: Overall view of *Trustball* installed in situ

Trustball is an arcade-style tangible game consisting of (a) a screen, (b) tangible LED buttons and (c) a container of custom-filled vending machine balls with a mechanism that would drop a ball into (d) a series of flippers that guided the ball in accordance with participants’ question responses (see in figure 2). Similar to the Voxbox design [31], this ball was aimed to incentivize participants to continue answering questions to gain a reward. The blinking buttons and start screen of *Trustball* invited visitors to “Press any button to play” – denoting the start of an interaction session. After pressing, the second page confronted the participant with terms and conditions outlining our actual research study’s terms and allowing them to proceed by pressing buttons to either “Agree” or “Disagree” with these terms. Both responses advanced the participant to the next screen, allowing all participants to play the game and answer questions with the only difference being data use in this study. The next screen confronted participants with the time it took them

to make their choice – acting as a deliberate provocation to encourage participants to think about their approach to the T&Cs. While this screen was shown to participants, a ball – metaphorically containing the ‘players data’ – was released into the game mechanism to roll and rest at the first flipper waiting for further interactions. The following three screens each showed a randomised scenario from a set of six privacy and data sharing scenarios followed by the question, **who would you trust to make a decision on your behalf?**, and the possible responses displayed on the screen above the 5 buttons. Each time a participant answered a question the data ball dropped into the next section. To release their data ball at the bottom of the machine (shown as (e) in figure 2) participants had to answer all three questions. The reward for this interaction was the ball itself, which initially contained a piece of candy, a data provocation relating to the scenarios and the relevant study information (project and researcher’s details). After discovering that users tended to repeat interactions to expedite the release of more candy, we removed the sweet and solely provided the participants with the provocation and information in the data ball.

Consent Scenarios and Delegates

In order to explore attitudes to consent delegation we needed a set of example scenarios illustrative of the broad range of concerns that may arise in typical T&C consent decisions and yet were sufficiently simple to be easily understood by participants from the general public.

Privacy and consent are widely recognised as complex, multifaceted issues. To ensure that our scenarios reflect this complexity we drew inspiration from both theoretical frameworks and real-world examples. Specifically, we were motivated by theoretical sets of privacy dimensions [14, 17, 20, 26, 70, 73] and in particular, by the proposal of Roessler [68] who argues that theories of privacy are simultaneously theories about protecting individual liberties and proposes three distinct dimensions of privacy, namely 1. *Decisional privacy* (privacy of actions and decisions, especially about people’s identity, way of life or projects they pursue without interference of others, self-determination); 2. *Informational privacy* (the perceived control of distribution of personal data, connected to notions of privacy violations through surveillance and protection of freedom, autonomy and agency of the individual); 3. *Local Privacy* (the concept that individuals rights and freedoms develop in solitude and are conditional on having a space (home) to withdraw to from public). Drawing on these dimensions of privacy we explored different categories of popular applications and constructed a set of scenarios (see Table 1) that were grounded in real-world examples, reflected the complexity highlighted by the theoretical frameworks, and captured the ambiguity and vagueness of contemporary consent decisions.

For each scenario participants could select from four potential delegates (tab. 2), derived from the literature on recommendation systems and technical feasibility. In addition to these delegation options, we added *Myself* as an option to enable participants to signal that they would not be prepared to delegate the decision and wished to retain control.

Festival Deployment

During the 2018 Edinburgh International Festival, we installed *Trustball* for 3 weeks alongside a series of interactive designs as part of a temporary exhibition entitled *Data Pipe Dreams: Glimpses of a Near Future*. Edinburgh International Festival is an annual cultural event attracting large numbers of local and international visitors across several weeks including numerous free public engagement activities, performances and entertainment areas. This exhibition was installed in a public pavilion in the city centre as part of one of the festival's entertainment areas and was located in the vicinity of other bars, food stalls, cinema and performance venues. The exhibition opened every day between 11am and 6pm with three evening events (opening, interim and closing) which lasted from 6pm until 9pm. The exhibition and *Trustball* installation were self-contained and visitors to the city could freely walk into and around the exhibition space.

As part of the *Trustball* installation, we captured a mix of quantitative and qualitative data. Firstly, we captured each interaction event and responses to the questions posed on the screen, alongside the response time and simple demographic information. Secondly, we conducted a set of follow-up questionnaires with participants who finished interacting with *Trustball*. And finally, we collected approximately 7 days of field observations to contextualise the interactions.

Data Pre-Processing and Cleansing

Over the entire duration of the festival, we counted 6,799 visitors in our temporary exhibition. While not all of these visitors engaged with every exhibit, we recorded a total of 1,749 unique interactions. Participants were predominantly from UK and European backgrounds ranging from university colleagues with technology experience to cultural event visitors, couples, families and general passers-by with little to no technology background. Due to the nature of the deployment in the context of a festival, we captured a large amount of noise in the *Trustball* interactions as a result of participants focusing on the gaming aspect of *Trustball* and selecting answers without reading or considering the scenarios.

In order to clean the dataset and filter valid participant responses, we specifically considered the reading times of participants for individual questions. Previous research suggests an average reading speed of 184 words per minute [76]. In the context of *Trustball*, the longest scenario consists of 53 and the shortest scenario of 20 words yielding an average

reading time of 6.5–17.2 seconds. To account for participants who read faster, we specified a 'minimum reading time' of 5 seconds for the short and 7 seconds for the long scenario, and filtered out all responses of participants who answered a scenario faster than the specified minimum reading time leaving us with a set of 565 participants (i.e. around 32% of the 1,749 interacting visitors). We confirmed the validity of the filter by plotting the distribution of response times – showing two peaks at approximately 1 and 10 seconds. Upon application of the filtering, the first peak of fast responses was removed and the density plots showed an even distribution of response times.

4 RESULTS AND ANALYSIS

We begin our analysis of the *Trustball* installation by describing observations captured by researchers in the field – thus setting the scene for the subsequent quantitative analysis of the *Trustball* dataset. To correlate participants' interactions with their perceptions of consent and delegation, we additionally provide an analysis of the follow-up questionnaires.

Participant Observations

Participants of *Trustball* fell into one of three categories, i.e. individual players (interacting with *Trustball* alone or without interference from others); pairs/couples (with both people standing in front of *Trustball*, but with varying degrees of interaction); and, groups of more than two people.

The display of the time participants had taken to read the *Trustball* T&Cs (see *Designing Trustball*) acted as a strong provocation – triggering reactions from all participants and prompting discussions among those participants in pairs or groups. The most common initial reaction from members of pairs or groups was laughter at “being busted” for not reading the T&Cs, often followed by admission that they never read T&Cs. This is contrary to peoples' claims in the questionnaires (see *Questionnaires*), but can perhaps be explained by adapting to social desirability and expected behaviour [41]. A similar reaction could be observed in individual participants, who often smiled, or took a slight step back, signalling surprise at the unexpected provocation. When in pairs, some participants started reading the T&Cs, but were pushed by their companions to “just click agree” who then. Interestingly, when the provocation appeared, these same companions then appeared critical: “See, this is what happens when you don't read it!”, “I told you to always read them!”.

When playing the game itself, couples exhibited a range of behaviours including reading the options out loud, discussing (or not discussing options), taking turns when pressing the buttons and, handing over part-completed games to their partner. Groups always discussed the options, but most of the time one person, usually the one who started it, was in charge of pressing the buttons. While groups in general

Table 1: Consent scenarios.

Roessler's Framework	Layers of Privacy	Scenario	Description
Informational Privacy	Mental Privacy	S1: Entertainment History	A game asks you to share music playlists and listening habits to develop an algorithm to change in-game-music according to your taste and mood.
Local Privacy	Spatial Privacy	S2: Location Data	Your weather app request access to your location and audio data to investigate noise pollution in your city. This data will be shared with the local council to review speed limits.
Local Privacy	Interactional Privacy	S3: Browsing History	Your sibling shares a survey with you to plan a joint holiday. They use an online service that asks for access to your entire browser history.
Informational Privacy	Informational Privacy	S4: Contact Lists	Your boss requires you to download a new chat service to communicate with your colleagues. You download the service and it asks for access to all your contacts in your address book, not only work-related ones. If you don't agree, you can't use the service for work.
Decisional Privacy	Decisional Privacy	S5: Social Media Activity	To provide personalised food boxes, local producers ask you to share social media likes of food pictures.
Decisional Privacy	Bodily Privacy	S6: Health Data	Public health services ask you share anonymised medical records with third parties to improve services.

Table 2: Potential delegation options.

Delegation	Description
<i>Myself</i>	The status quo of keeping control of user agreements and continuing to agree to T&C on a one click basis.
<i>Friend</i>	Selecting a trusted friend to make decisions on one's behalf.
<i>Expert</i>	Not further specified, this could include a legal authority, lawyer or IT/data experts.
<i>Crowd</i>	Basing decision on the majority of users of a service rather than an individual choice.
<i>AI/Bot</i>	An AI concept learning data sharing preferences to become a more automated consent response.

had the most lively discussions, differences in discussion style could be observed with groups that appeared to be families or peers. In peer groups, discussions seemed to be more open, but they didn't always reach consensus. In families, two distinct behaviours showed to be prevalent: First, younger members sought guidance based on their own confidence towards the topic of the questions. The question that was most confusing to them related to workplace privacy where they often sought guidance from their older family members. We suspect this is due to their limited workplace experience. Second, when the main participant of the game was confident (and quick) in proclaiming the answer before actually clicking, sometimes other members would contest and give reasons to challenge their initial answer, e.g. (when discussing S6: Health Data) ... "What? No, in those things I only trust myself." (P1), "Really, don't you trust what your doctor says?" (P2), "Yeah, I wouldn't know if I always know best, I'd ask someone." (P3), "I guess you are right. Maybe an expert

then." (P1). In this example the dialogue led the participant to rethink their answer, demonstrating the influence of friends and families in making such decisions.

Trustball Interactions

Are Participants Prepared to Delegate Decision Making?

Our initial analysis of the interaction data focused on determining whether participants were prepared to delegate decision making in any of our scenarios.

Figure 3: Proportion of responses selecting *Myself* compared to any other delegation option for each scenario.

As shown in figure 3, the overall popularity of delegation (50.4% of responses) is approximately equal to the desire to retain control (49.6%). We observe significant differences across scenarios (min. 37% for S3: Browsing History and max. 61% for S1: Entertainment History). In particular, S3: Browsing History and S4: Contact Lists show a clear tendency for participants to opt not to delegate decisions, whilst participants appeared to be less opposed to delegation for S6: Health Data, S5: Social Media Activity and S1: Entertainment History. We observed a very subtle preference towards *Myself* for S2: Location Data. Using an uncorrected N-1 Chi-squared

test to conduct a pairwise comparison between scenarios, we can show significantly more participants selected *Myself* for S1: Entertainment History than S3: Browsing History ($p < 0.0001$) whilst for other comparisons (e.g. S5: Social Media Activity and S6: Health Data) we cannot show significance. Note, however, that our intention was to conduct a first exploration into willingness of participants to delegate consent rather than directed testing of a hypothesis – in keeping with the suggestions of [18] we therefore refrained from conducting an extensive significance analysis.

Does the Choice of Delegation Vary by Scenario? While many participants expressed willingness to delegate decision making, the choice of who to delegate to appears to be highly dependent on the scenario. Figure 4 provides an overview of the proportions of chosen delegation options per scenario. For example, *Friend* is the most popular delegation option (excluding *Myself*) for S1: Entertainment History, S3: Browsing History, and S5: Social Media Activity privacy scenarios whilst participants prefer an *Expert* for S2: Location Data, S4: Contact Lists and S6: Health Data. Both *AI/Bot* and *Crowd* appear to be consistently the least popular delegation option across all scenarios.

Figure 4: Proportion of responses selecting each delegation option (excluding *Myself*) for each scenario.

Impact of Demographics. Approximately 56% (250/565) of participants provided demographic data regarding their age group. We see roughly even distribution across the three youngest age brackets (16-25: 17.35%, 26-40: 21.42%, 41-60: 14.16%); just 2.83% of participants have a reported age of over 61 years. One might have expected delegation choices to differ by demographic. However, we find delegation preferences to be similar across all age groups (fig. 5).

Do Participants Make Consistent Delegation Choices? We were particularly interested in exploring whether individual participants had a tendency towards a specific delegation option (or chose not to delegate and retain control). Figure 6

Figure 5: Percentage of delegation choices by scenario and age group (total n=565).

Figure 6: Proportions of responses selecting one, two and three identical delegation options.

Figure 7: Proportion of responses selecting the delegation option at least one, two and three times.

Figure 8: Cumulative proportion of responses selecting delegation options zero, one, two and three times.

shows the proportion of participants selecting one, two and three different delegation options within their answer set. The data shows a clear tendency for participants to include two delegation options within their answers (257 participants) versus selecting the same delegation option for all three scenarios (156) or selecting a different delegation option for each scenario (152).

To further understand the variance in delegation choice exhibited by participants we show the proportion of participants selecting the same delegation option at least one, two and three times (fig. 7) and provide a plot of the cumulative distribution of responses (fig. 8). Considering *Myself* as a delegation option, over 70% of participants selected it at least once, $\approx 50\%$ at least twice and over 20% three times – confirming *Myself* as the most popular delegation option across all participants. Considering the remaining set of delegation options, we observe two clusters with similar popularity patterns: (1) *Friend* and *Expert*, and (2) *AI/Bot* and *Crowd*. Less

than 50% of participants selected any of the delegation options (except *Myself*) at least once; whilst cluster (1) was the most popular among the remaining set of delegation options ($\approx 40\%$) that were selected at least once. We note that *AI/Bot* and *Crowd* delegation options were not selected three times by any participant – emphasising the variance of delegation choices based on scenarios and highlighting the importance of providing multiple delegation options to participants.

Do Participants' Decisions to Delegate Match Their Actions?

Our study required participants to read and agree to a set of T&Cs prior to playing each *Trustball* game. The time this took each participant provides some insight into the attention they gave the task. Overall, we observe that participants spent 9.88 seconds (SD: 9.43, mdn: 6.53) reading and agreeing to the displayed T&Cs. This is a contrast to the expected duration: the T&Cs are formed of 243 words yielding a minimum reading time of over 60 seconds – more than six times higher than the mean reading time. We further analysed the preferred choices of *Myself* as compared to other delegation options based on the time participants spent reading the initial T&Cs (fig. 9). We observe that participants who choose *Myself* twice and three times (on average) spent longer reading the initial T&Cs than participants who selected *Myself* at most once. In particular, participants who never chose *Myself* as a delegation option are characterised by quick reading times of T&Cs (mean: 9.02 seconds, SD: 10.28, mdn: 5.45).

While participants appeared to skim our T&Cs, we observe that participants read the individual scenarios within *Trustball* more carefully with a mean reading time of 17.49 seconds (SD: 7.45, mdn: 16.26) on page 1, 13.49 seconds (SD: 6.43, mdn: 11.90) on page 2, and 11.88 seconds (SD: 6.07, mdn: 10.60) on page 3. Reading times noticeably decrease as the game progresses which we attribute to participants becoming more familiar with the operation of *Trustball*.

Figure 9: Time taken to read T&Cs for participants who chose *Myself* as a delegation option zero, one, two and three times. Box plot shows median, first and third quartiles, with whiskers stretching to 1.5 times the interquartile range. Data outside of this range is represented as outliers.

Questionnaires

After the game, participants were recruited at random to fill in a paper questionnaire. This survey contained qualitative and quantitative questions, and was not mapped to specific datasets but used to contextualise the *Trustball* interactions and resulting data overall. Our aim was to capture further insights into attitudes towards T&Cs and the concept of delegating consent decisions. In particular, we sought to understand how often participants actually read T&Cs they agree to (and how they make this decision), whether or not they had been asked by others to provide advice regarding T&Cs agreements, and who they would delegate consent decisions to. We collected 44 valid questionnaires.

Participant Practices when Agreeing to T&Cs. The majority of participants indicated that they rarely read T&Cs before agreeing to them (never: 40.9%, sometimes: 40.9%, frequently: 13.6%, all the time: 4.5%) confirming our observations from *Trustball* interactions in which the reading times for the T&Cs were less than the reading times of any question. When deciding whether to read T&Cs participants claimed to consider three factors: data importance, data recipients and data usage. Firstly, participants appear to base their decision (whether or not to read associated T&Cs) on the data the service is processing. For example, some participants claim to read T&Cs if financial transaction data is used. Secondly, participants claimed to consider not reading T&Cs related to services they trust (not reading T&Cs “[...] unless something very important or company felt I couldn’t trust [sic]”). Participants named “reputation”, “recognisability” and “established companies” as factors for their trust. Secondly, participants judged whether to accept T&Cs (and consequently consent to data usage and sharing) based on the legitimacy and plausibility of the service requiring specific data. In particular, participants expressed fear of improper data use, the resulting risks of becoming a victim of a hacker and the implications of public data exposure, or economical exploitation. In contrast, participants mentioned “health data for research” as positive examples of data usage ultimately benefiting the “greater good”.

Delegation Preferences and Practices. To further understand the rationale behind participants’ delegation choices we asked questions relating to when, why and to whom they might choose to delegate consent. Making consent decisions themselves was viewed as important to many participants, even if they sought advice in making this decision, as they wanted to “have the last say” and maintain control. Generally, participants reasoned that they trust their own judgement, instincts and knowledge from past experiences. However, some participants acknowledged their own biases, their subjective judgement and contextualised that decisions are made

on an everyday basis, which require reevaluation *in situ*. In a few rare occasions participants claimed not to think, mind or care about their T&Cs decisions.

We asked participants to rate on a scale of 1-9 (1: not useful at all; 9: very useful) the ‘usefulness’ of delegation options in situations when delegating decision making was appropriate. Based on these ratings we created a ranking of most and least useful delegation options: *Expert* (54%), *Friend* (31%), *AI/Bot* (10%) and *Crowd* (6%). In addition to the quantitative rating, we further asked participants to explain their choices. Both the ranking of the delegation options and the qualitative analysis of responses highlight that friends and expert were the most popular choices – similar to the *Trustball* results. In particular, participants mentioned friends, partners and family, and in some rare cases colleagues. The main criteria for these choices was expressed as the closeness or intimacy of the relationship to the person to whom they were willing to delegate decisions. This criteria was connected with beliefs that delegates would have good intentions and hence “act in best interest” and that due to the closeness, they would make similar and/or acceptable decisions (“know me well”, “like-minded”) on the participants’ behalf. While some participants recognised that “even friends have agendas” and close persons still have biases, others would see their close ties as more wise (“Mum always knows best”) or more qualified to make decisions than themselves.

Experts were viewed as domain-specific entities (e.g. financial or health domain) and as legitimate and trustworthy delegates. While an *Expert* is trusted by default, *AI/Bot* and the more general *Crowd* appear to have to gain trust to be viewed as a source of credible information and hence legitimate advisory in data decisions – the ranking shows the current very low popularity of these delegation options. Factors leading to trust identified by participants were (1) the entity (e.g. company or institution) and their partners who developed and programmed (“coded”) the AI, (2) on the way how it was trained and with what/whose kind of data, and (3) that the AI is proven to work correctly.

Finally, we explored the extent to which participants had acted as T&Cs ‘delegates’ or ‘recommenders’ for others. For many this appeared to be a common occurrence (‘all the time’: 7.0%, ‘frequently’: 32.6%, ‘sometimes’: 53.5%, ‘never’: 7.0%) and provides further evidence that reframing consent in terms of delegation aligns with common practices.

5 DISCUSSION

In seeking to understand consent decision-making and its potential reframing as acts of delegation, our initial objective was to explore the extent to which participants would be prepared to delegate consent decisions. Our results show that approximately 50% of participants were accepting of delegation options and suggest that this reframing warrants

further exploration by the research community given the shortcomings of existing consent approaches.

Understanding Delegation Choices

Preliminary insights into participants’ choices of *who* to delegate to and their stated rationale highlight the importance of trust and perceived control over the outcomes of the consent decision. The three distinct factors contributing to trust in a delegatee were how the participant perceived 1. their competence or abilities to make authoritative decisions; 2. their intention in the decision-making process and 3. their moral integrity. The actual selection of a delegatee was based either on an already established trust relationship or a dependence, e.g. due to the lack of own knowledge in the field. The most popular choices for delegatee were those with close personal ties such as friends/family members (who possess good intentions or moral integrity) or experts (who are perceived as knowledgeable). Delegation in this sense meant that the delegator expected the outcome to be in a range of “acceptable” options and hence predictable to a certain degree. As a result of these stable expectations, one can speculate that delegation within trust relationships leads to a reduction of complexity and uncertainty hence to an increased feeling of security. Simply put, even if people delegate, they still feel like they are in control.

Those who did not delegate consent decisions rationalised their choice of “myself” in terms of their desire for control and mistrust of either the data controllers or the delegation options offered. In particular, participants often cited concern that the risks and consequences of what they considered as a “wrong” decision outweighed the perceived advantages of delegation. Delegation for these participants appeared to be perceived as a loss of control. Equally, for less popular delegation options, such as *Crowd* and *AI/Bot*, the lack of trust played a part. The *Crowd* delegatee, for example, seemed ill-defined to participants and therefore it was hard for them to assess its trustworthiness. This could have resulted from our choice of the word, for example the term ‘community’ might have potentially provided different results. It is, however, interesting to note that while not many of our participants appeared to trust *Crowd*, many probably rely on crowd-sourced recommendations for their holidays (e.g. Trip Advisor [78]), financial decisions (e.g. MoneySupermarket [54]), employees or collaborators (e.g. LinkedIn [46]) and so on. In the case of *AI/Bot*, it was more evident what conditions lead to a trust relationship. Participants viewed *AI/Bot* as potentially unfair and biased. In order to become a valid delegation option, participants wanted reassurance that the development of *AI/Bot* was controlled and regulated to reduce the risks of unwanted outcomes. As this evaluation can’t be done by the individual alone, the question of control and regulation becomes a collective matter. In summary, our

results provided insights into if and under what conditions participants choose delegation and highlight the importance of a sense of agency, trust and risk for participants in making these decisions.

Reflecting on Public Engagement

The results from our study have shown that *Trustball* was successful in engaging public audiences with questions around consent and delegation in a tangible form. The use of a tangible questionnaire as a study method was labor-intensive but allowed for richer exploration than, e.g. purely online surveys. While there is always “an element of risk in deploying probes; they might fail or bring unexpected results” [37, p. 18], we found a strong methodology to analyse the data generated from this probe was required to validate the data, as discussed by Boehner et al. [13]. While an online survey would have potentially collected more individual responses, the social context and interactions of *Trustball* seemed to mirror many of the pressures present when making traditional consent decisions – participants were regularly confronted with distractions, time and peer-pressures by either friends or family they were with, other planned activities fighting for their attention. If one expects and accepts the individual choices of user agreements, these more complex, contextual and social situations in which decisions are made are often undervalued. We believe this aspect should remain part and feature more heavily in research of data capture processes.

Limitations of our Study

There are several limitations to our study and the extent to which we were able to analyse and interpret the captured data. Firstly, the nature of the in-the-wild deployment in the context of this festival exhibition led to significant noise in the captured dataset – requiring potentially error-prone filtering of invalid events based on low reading and interaction durations. In a small number of cases we observed participants being replaced by other members of their group while a game was in progress – an unexpected behaviour that cannot easily be recognised from the captured interaction data and could influence our findings when we aggregate multiple questions to report on participant preferences. However, we believe the number of occurrences of this behaviour to be small. Most significantly, our choices of scenarios and delegation options clearly influence participant behaviour. However, given that a willingness to delegate was visible in all of the scenarios we believe that it is likely that our results are robust to the introduction of new scenarios. While we purposefully kept a forced choice for clear delegation responses, we did not intend to conflate complex consent and advice behaviours and acknowledge further differentiation between these concepts in decision-making processes.

6 CONCLUDING REMARKS – CONSENT BEYOND THE INDIVIDUAL

In this paper, we have presented an interactive questionnaire as technology probe to investigate how the public would respond to the concept of delegating consent. Although there are definite limitations to this study, we have identified clear tendencies and public interest which warrants further investigation into alternative systems for privacy management, consent and delegation for HCI researchers. While some research within this area is focusing on locking down user’s personal data [16] for protection of privacy, other scholars recommend design guidelines to support HCI researchers as well as commercial developers in the design of informed and meaningful consent mechanisms or systems. For example, Lugger and Rodden suggest systems that allow for data review and consent withdrawal at any time, scaffold understanding through visualisations and opportunities for further interrogation [49, 50]. Equally, they emphasise the “need for multiple approaches if effective consent practices are to be realised within the design of pervasive systems” [50, p.393]. Our paper contributes to these areas of debate and the series of alternative consent mechanisms or systems with the proposition to open consent out instead of closing it down. It has increasingly been argued that informed consent is paradoxical and that we need to move beyond the individual “click-wrap agreement culture” [6, p. 22] to alternative models of consent [19, 50].

Our study confirms these shortcomings of the existing models and provides us with a series of considerations relevant for the design of future consent systems. The varied results we encountered for the scenarios highlights the necessity for flexible, adaptable and personalisable delegation mechanisms which allow the user to control not only *if* they consent but *how* they choose to consent. As shown, in some instances users may choose to read the ‘fine print’ carefully while in other situations they may accept an automated or delegated decision more comfortably. By offering such a flexible system, the user has the opportunity to make situation-specific decisions with each new data request of when to control a data agreement, when to trust others or when to automate consent. As identified by our participants, this may depend on factors such as data usage, data recipients or data importance but relationships to delegates played at least an equally important role in participant’s willingness to trust a consent decision to a potential delegatee. We therefore propose that in order to offer users control beyond their current perceptions of being in charge, real agency lies in the offer of choice. So we emphasise the importance of multiple options for self-declaration, delegation and beyond in flexible systems for consent. Our future work will investigate such a

system in more detail to explore how an overarching flexible infrastructure may offer users more choice, control and ultimately agency and give them the tools to make informed choices about when to consent, when to delegate and when to automate decision-making. Reframing models of consent beyond the individual reconsiders protection of (individual) privacy for a wider collective democratic participation in digital environments.

7 ACKNOWLEDGEMENTS

This research is partially funded by the UK EPSRC under grant number EP/N028228/1 (PACTMAN) and by the Department of Health as part of the Connected Health Cities project. We thank Siyao Zhou and Joanna Spreadbury for helping make Trustball.

REFERENCES

- [1] [n. d.]. Privacy-aware Location Privacy Preference Recommendations (*MOBIQUITOUS '14*). ICST, Brussels, Belgium. <https://doi.org/10.4108/icst.mobiquitous.2014.258017>
- [2] Elena Agapie, Lucas Colusso, Sean A. Munson, and Gary Hsieh. 2016. PlanSourcing: Generating Behavior Change Plans with Friends and Crowds. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, New York, NY, USA, 119–133. <https://doi.org/10.1145/2818048.2819943>
- [3] Yuvraj Agarwal and Malcolm Hall. 2013. ProtectMyPrivacy: detecting and mitigating privacy leaks on iOS devices using crowdsourcing. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*. ACM, Taipei, Taiwan, 97–110. <https://doi.org/10.1145/2462456.2464460>
- [4] Tim Baarslag, Alper T. Alan, Richard Gomer, Muddasser Alam, Charith Perera, and Enrico H. Gerding. 2017. An Automated Negotiation Agent for Permission Management. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, São Paulo, Brazil, 380–390.
- [5] Solon Barocas and Helen Nissenbaum. 2014. Big Data's End Run Around Procedural Privacy Protections. *Commun. ACM* 57, 11 (Oct. 2014), 31–33. <https://doi.org/10.1145/2668897>
- [6] Anja Bechmann. 2014. Non-Informed Consent Cultures: Privacy Policies and App Contracts on Facebook. *Journal of Media Business Studies* 11, 1 (March 2014), 21–38. <https://doi.org/10.1080/16522354.2014.11073574>
- [7] Rajesh Bejugam and Kirsten LeFevre. 2011. enList: Automatically Simplifying Privacy Policies. In *2011 IEEE 11th International Conference on Data Mining Workshops*. IEEE, Vancouver, BC, Canada, 620–627. <https://doi.org/10.1109/ICDMW.2011.74>
- [8] Ardion Beldad, Thea van der Geest, Menno de Jong, and Michael Steehouder. 2012. Shall I Tell You Where I Live and Who I Am? Factors Influencing the Behavioral Intention to Disclose Personal Data for Online Government Transactions. *International Journal of Human-Computer Interaction* 28, 3 (March 2012), 163–177. <https://doi.org/10.1080/10447318.2011.572331>
- [9] Ruha Benjamin. 2016. Informed Refusal: Toward a Justice-based Bioethics. *Science, Technology, & Human Values* 41, 6 (Nov. 2016), 967–990. <https://doi.org/10.1177/0162243916656059>
- [10] Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, and Tom Yeh. 2010. VizWiz: Nearly Real-time Answers to Visual Questions. In *Proceedings of the 2010 International Cross Disciplinary Conference on Web Accessibility (W4A '10)*. ACM, New York, NY, USA, Article 24, 2 pages. <https://doi.org/10.1145/1805986.1806020>
- [11] Igor Bilogrevic, Kévin Huguenin, Berker Agir, Murtuza Jadhliwala, Maria Gazaki, and Jean-Pierre Hubaux. 2016. A machine-learning based approach to privacy-aware information-sharing in mobile social networks. *Pervasive and Mobile Computing* 25 (Jan. 2016), 125–142. <https://doi.org/10.1016/j.pmcj.2015.01.006>
- [12] Igor Bilogrevic, Kévin Huguenin, Murtuza Jadhliwala, Florent Lopez, Jean-Pierre Hubaux, Philip Ginzboorg, and Valtteri Niemi. 2013. Inferring social ties in pervasive networks: an on-campus comparative study. In *Proceedings of the 2013 ACM conference on Pervasive and ubiquitous computing adjunct publication - UbiComp '13 Adjunct*. ACM Press, Zurich, Switzerland, 123–126. <https://doi.org/10.1145/2494091.2494128>
- [13] Kirsten Boehner, Janet Vertesi, Phoebe Sengers, and Paul Dourish. 2007. How HCI Interprets the Probes. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 1077–1086. <https://doi.org/10.1145/1240624.1240789>
- [14] Judee K. Burgoon. 1982. Privacy and Communication. *Annals of the International Communication Association* 6, 1 (Jan. 1982), 206–249. <https://doi.org/10.1080/23808985.1982.11678499>
- [15] France Bélanger, Robert E. Crossler, Janine S. Hiller, Jung-Min Park, and Michael S. Hsiao. 2013. POCKET: A tool for protecting children's privacy online. *Decision Support Systems* 54, 2 (Jan. 2013), 1161–1173. <https://doi.org/10.1016/j.dss.2012.11.010>
- [16] Amir Chaudhry, Jon Crowcroft, Heidi Howard, Anil Madhavapeddy, Richard Mortier, Hamed Haddadi, and Derek McAuley. 2015. Personal Data: Thinking Inside the Box. *Aarhus Series on Human Centered Computing* 1, 1 (Oct. 2015), 4. <https://doi.org/10.7146/aahcc.v1i1.21312>
- [17] Hyunyi Cho and Robert Larose. 1999. Privacy Issues in Internet Surveys. *Social Science Computer Review* 17, 4 (Nov. 1999), 421–434. <https://doi.org/10.1177/089443939901700402>
- [18] Andy Cockburn, Carl Gutwin, and Alan Dix. 2018. HARK No More: On the Preregistration of CHI Experiments. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, 141:1–141:12. <https://doi.org/10.1145/3173574.3173715>
- [19] Bart Custers. 2016. Click here to consent forever: Expiry dates for informed consent. *Big Data & Society* 3, 1 (Jan. 2016), 2053951715624935. <https://doi.org/10.1177/2053951715624935>
- [20] Tobias Dienlin. 2014. The Privacy Process Model. In *Medien und Privatheit*, M. Herz & J.-M. Mönig S. Garnett, S. Halft (Ed.). Stutz, Passau, Germany, 105–122.
- [21] Djellel Difallah, Elena Filatova, and Panos Ipeirotis. 2018. Demographics and Dynamics of Mechanical Turk Workers. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining (WSDM '18)*. ACM, New York, NY, USA, 135–143. <https://doi.org/10.1145/3159652.3159661>
- [22] Andy Dow, John Vines, Rob Comber, and Rob Wilson. 2016. ThoughtCloud: Exploring the Role of Feedback Technologies in Care Organisations. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 3625–3636. <https://doi.org/10.1145/2858036.2858105>
- [23] Thomas W. Dunfee and Thomas Donaldson. 1995. Contractarian Business Ethics: Current Status and Next Steps. *Business Ethics Quarterly* 5, 2 (1995), 173–186. <https://doi.org/10.2307/3857352>
- [24] Ulrike Felt, Milena D. Bister, Michael Strassnig, and Ursula Wagner. 2009. Refusing the Information Paradigm: Informed Consent, Medical Research, and Patient Participation. *Health: An Interdisciplinary Journal for the Social Study of Health, Illness and Medicine* 13, 1 (2009), 87–106. <http://hea.sagepub.com/content/13/1/87.short> 00053.

- [25] Holly Fernandez Lynch, Steven Joffe, and Eric A. Feldman. 2018. Informed Consent and the Role of the Treating Physician. *New England Journal of Medicine* 378, 25 (June 2018), 2433–2438. <https://doi.org/10.1056/NEJMHle1800071>
- [26] Luciano Floridi. 1999. Information ethics: On the philosophical foundation of computer ethics. *Ethics and Information Technology* 1, 1 (March 1999), 33–52. <https://doi.org/10.1023/A:1010018611096>
- [27] Ricard L. Fogues, Pradeep K. Murukannaiah, Jose M. Such, and Munindar P. Singh. 2017. SoSharP: Recommending Sharing Policies in Multiuser Privacy Scenarios. *IEEE Internet Computing* 21, 6 (Nov. 2017), 28–36. <https://doi.org/10.1109/MIC.2017.4180836>
- [28] Sarah Gallacher, Connie Golsteijn, Yvonne Rogers, Licia Capra, and Sophie Eustace. 2016. SmallTalk: Using Tangible Interactions to Gather Feedback from Children. In *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '16)*. ACM, New York, NY, USA, 253–261. <https://doi.org/10.1145/2839462.2839481>
- [29] Sarah Gallacher, Connie Golsteijn, Lorna Wall, Lisa Koeman, Sami Andberg, Licia Capra, and Yvonne Rogers. 2015. Getting Quizzical About Physical: Observing Experiences with a Tangible Questionnaire. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 263–273. <https://doi.org/10.1145/2750858.2807529>
- [30] Bill Gaver, Tony Dunne, and Elena Pacenti. 1999. Design: Cultural Probes. *interactions* 6, 1 (Jan. 1999), 21–29. <https://doi.org/10.1145/291224.291235>
- [31] Connie Golsteijn, Sarah Gallacher, Lisa Koeman, Lorna Wall, Sami Andberg, Yvonne Rogers, and Licia Capra. 2015. VoxBox: A Tangible Machine That Gathers Opinions from the Public at Events. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '15)*. ACM, New York, NY, USA, 201–208. <https://doi.org/10.1145/2677199.2680588>
- [32] Pauline Gourlet and Thierry Dassié. 2017. Cairn: A Tangible Apparatus for Situated Data Collection, Visualization and Analysis. In *Proceedings of the 2017 Conference on Designing Interactive Systems (DIS '17)*. ACM, New York, NY, USA, 247–258. <https://doi.org/10.1145/3064663.3064794>
- [33] Mats G Hansson, Joakim Dillner, Claus R Bartram, Joyce A Carlson, and Gert Helgesson. 2006. Should donors be allowed to give broad consent to future biobank research? *The Lancet Oncology* 7, 3 (March 2006), 266–269. [https://doi.org/10.1016/S1470-2045\(06\)70618-0](https://doi.org/10.1016/S1470-2045(06)70618-0)
- [34] Yoan Hermstrüwer and Stephan Dickert. 2017. Sharing is daring: An experiment on consent, chilling effects and a salient privacy nudge. *International Review of Law and Economics* 51, Supplement C (Aug. 2017), 38–49. <https://doi.org/10.1016/j.irle.2017.06.001>
- [35] Hwajung Hong, Eric Gilbert, Gregory D. Abowd, and Rosa I. Arriaga. 2015. In-group Questions and Out-group Answers: Crowdsourcing Daily Living Advice for Individuals with Autism. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 777–786. <https://doi.org/10.1145/2702123.2702402>
- [36] Steven Houben and Christian Weichel. 2013. Overcoming Interaction Blindness Through Curiosity Objects. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13)*. ACM, New York, NY, USA, 1539–1544. <https://doi.org/10.1145/2468356.2468631>
- [37] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B. Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, Nicolas Roussel, and Björn Eiderbäck. 2003. Technology Probes: Inspiring Design for and with Families. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*. ACM, New York, NY, USA, 17–24. <https://doi.org/10.1145/642611.642616>
- [38] Randall Kiser. 2010. *Beyond Right and Wrong: The Power of Effective Decision Making for Attorneys and Clients*. Springer Science & Business Media, Berlin-Heidelberg, Germany. Google-Books-ID: xEUP530LI2MC.
- [39] Bart Knijnenburg and Alfred Kobsa. 2013. Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 3, 20 (Oct. 2013), 23. <https://doi.org/10.1145/2499670>
- [40] Sherrie Y. X. Komiak and Izak Benbasat. 2006. The Effects of Personalization and Familiarity on Trust and Adoption of Recommendation Agents. *MIS Quarterly* 30, 4 (2006), 941–960. <http://www.jstor.org/stable/25148760>
- [41] Tuukka Lehtiniemi and Yki Kortensniemi. 2017. Can the obstacles to privacy self-management be overcome? Exploring the consent intermediary approach. *Big Data & Society* 4, 2 (Dec. 2017), 2053951717721935. <https://doi.org/10.1177/2053951717721935>
- [42] Lei Li, Tong Sun, and Tao Li. 2011. Personal Social Screen—A Dynamic Privacy Assignment System for Social Sharing in Complex Social Object Networks. IEEE, Boston, MA, USA, 1403–1408. <https://doi.org/10.1109/PASSAT/SocialCom.2011.147>
- [43] Qingrui Li, Juan Li, Hui Wang, and Ashok Ginjala. 2011. Semantics-enhanced privacy recommendation for social networking sites. In *Trust, Security and Privacy in Computing and Communications (TrustCom)*, 2011 IEEE 10th International Conference on. IEEE, Ottawa, ON, Canada, 226–233.
- [44] Christian Fernando Libaque-Saenz, Younghoon Chang, Jimin Kim, Myeong-Cheol Park, and Jae Jeung Rho. 2016. The role of perceived information practices on consumers' intention to authorise secondary use of personal data. *Behaviour & Information Technology* 35, 5 (May 2016), 339–356. <https://doi.org/10.1080/0144929X.2015.1128973>
- [45] Jialiu Lin, Shahriyar Amini, Jason I. Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. 2012. Expectation and Purpose: Understanding Users' Mental Models of Mobile App Privacy Through Crowdsourcing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. ACM, New York, NY, USA, 501–510. <https://doi.org/10.1145/2370216.2370290>
- [46] LinkedIn. 2019. Retrieved from <https://www.linkedin.com/>. (2019).
- [47] Ewa Luger, Marina Jirotko, Tom Rodden, and Lilian Edwards. 2014. How do you solve a problem like consent?: the workshop. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM Press, New York, USA, 613–619. <https://doi.org/10.1145/2638728.2641676>
- [48] Ewa Luger, Stuart Moran, and Tom Rodden. 2013. Consent for All: Revealing the Hidden Complexity of Terms and Conditions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2687–2696. <https://doi.org/10.1145/2470654.2481371>
- [49] Ewa Luger and Tom Rodden. 2013. An Informed View on Consent for UbiComp. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. ACM, New York, NY, USA, 529–538. <https://doi.org/10.1145/2493432.2493446>
- [50] Ewa Luger and Tom Rodden. 2014. The value of consent: Discussions with designers of ubiquitous computing systems. In *2014 IEEE International Conference on Pervasive Computing and Communication Workshops (PERCOM WORKSHOPS)*. IEEE, Zurich, Switzerland, 388–393. <https://doi.org/10.1109/PerComW.2014.6815237>
- [51] Tobias Matzner. 2014. Why privacy is not enough privacy in the context of "ubiquitous computing" and "big data". *Journal of Information Communication and Ethics in Society* 12 (06 05 2014), 93–106. Issue 2. <https://doi.org/10.1108/JICES-08-2013-0030>
- [52] Tobias Matzner and Carsten Ochs. 2017. Sorting Things Out Ethically: Privacy as a Research Issue beyond the Individual. In *Internet Research Ethics for the Social Age: New Challenges, Cases, and Contexts*, Michael

- Zimmer and Katherina Kinder-Kurlanda (Eds.). Peter Lang Publishing, New York, USA, 39–52.
- [53] Aleecia M. McDonald and Lorrie Faith Cranor. 2008. The Cost of Reading Privacy Policies. *LS: A Journal of Law and Policy for the Information Society* 4 (2008), 22. Issue 3.
- [54] Moneysupermarket. 2013. Retrieved from <http://www.moneysupermarket.com>. (2013).
- [55] Alistair Morrison, Donald McMillan, and Matthew Chalmers. 2014. Improving Consent in Large Scale Mobile HCI Through Personalised Representations of Data. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction (NordCHI '14)*. ACM, New York, NY, USA, 471–480. <https://doi.org/10.1145/2639189.2639239>
- [56] Helen Nissenbaum. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford Law Books, Stanford, CA.
- [57] Helen Nissenbaum. 2011. A Contextual Approach to Privacy Online. *Dædalus* 4 (2011), 32–48.
- [58] Chris Norval and Tristan Henderson. 2017. Contextual Consent: Ethical Mining of Social Media for Health Research. *arXiv:1701.07765 [cs]* (Jan. 2017). <http://arxiv.org/abs/1701.07765> arXiv: 1701.07765.
- [59] Jonathan A. Obar and Anne Oeldorf-Hirsch. 2016. *The Biggest Lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services*. SSRN Scholarly Paper ID 2757465. Social Science Research Network, Rochester, NY. <https://papers.ssrn.com/abstract=2757465>
- [60] Carsten Ochs. 2015. "Selbstschutz", oder:: Kollektive Privatheitspraktiken als politisches Handeln in digitalen Öffentlichkeiten. *Forschungsjournal Soziale Bewegungen* 3 (Nov. 2015), 45–54. https://www.wiso-net.de/document/FJSB_3CE4A11C6791F0975576100DFCB5A6E9
- [61] Carsten Ochs and Jörn Lamla. 2017. Demokratische Privacy by Design. *Forschungsjournal Soziale Bewegungen* 30, 2 (Jan. 2017), 189–199. <https://doi.org/10.1515/fjsb-2017-0040>
- [62] Carsten Ochs, Fabian Pittroff, Barbara Büttner, and Jörn Lamla. 2016. Governing the internet in the privacy arena. *Internet Policy Review* 5, 3 (Sept. 2016). <https://policyreview.info/articles/analysis/governing-internet-privacy-arena>
- [63] Terms of Service; Didn't Read. 2018. Retrieved from <https://tosdr.org/>. (2018).
- [64] Publications Office of the European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). Retrieved from <https://eur-lex.europa.eu/eli/reg/2016/679/oj>. (2016).
- [65] PersonalData.IO. 2018. Retrieved from <https://personaldata.io/>. (2018).
- [66] Bahman Rashidi, Carol Fung, and Tam Vu. 2015. Dude, ask the experts!: Android resource access permission recommendation with RecDroid. In *IFIP/IEEE International Symposium on Integrated Network Management (IM)*. IEEE, Ottawa, ON, 296–304. <https://doi.org/10.1109/INM.2015.7140304>
- [67] Iyubanit Rodríguez, Maria Karyda, Andrés Lucero, and Valeria Herzkovic. 2018. Exploring Tangible Ways to Evaluate User Experience for Elders. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18)*. ACM, New York, NY, USA, Article LBW589, 6 pages. <https://doi.org/10.1145/3170427.3188450>
- [68] Beate Roessler. 2006. New Ways of Thinking About Privacy. In *Oxford Handbook of Political Theory*, Anne Phillips Bonnie Honig and John Dryzek (Eds.). Oxford University Press, Oxford, 694–713.
- [69] Beate Roessler and Dorota Mokrosinska (Eds.). 2015. *Social Dimensions of Privacy*. Cambridge University Press, Cambridge. <https://doi.org/10.1017/CBO9781107280557.001>
- [70] Daniel J. Solove. 2002. Conceptualizing Privacy. *California Law Review* 90 (2002), 1087–1156. <https://heinonline.org/HOL/P?h=hein.journals/calr90&i=1101>
- [71] Daniel J. Solove. 2013. Introduction: Privacy Self-Management and the Consent Dilemma. *Harvard Law Review* 126 (2013), 1880–1903.
- [72] Anna C. Squicciarini, Elisa Bertino, Elena Ferrari, and Indrakshi Ray. 2006. Achieving privacy in trust negotiations with an ontology-based approach. *IEEE Transactions on Dependable and Secure Computing* 3, 1 (Jan. 2006), 13–30. <https://doi.org/10.1109/TDSC.2006.3>
- [73] Herman T. Tavani. 2008. Floridi's ontological theory of informational privacy: Some implications and challenges. *Ethics and Information Technology* 10, 2-3 (Sept. 2008), 155–166. <https://doi.org/10.1007/s10676-008-9154-x>
- [74] Nick Taylor, Justin Marshall, Alicia Blum-Ross, John Mills, Jon Rogers, Paul Egglesstone, David M. Frohlich, Peter Wright, and Patrick Olivier. 2012. Viewpoint: Empowering Communities with Situated Voting Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 1361–1370. <https://doi.org/10.1145/2207676.2208594>
- [75] Nicholas K. Taylor, Elizabeth Papadopoulou, Sarah Gallacher, and Howard M. Williams. 2013. Is There Really a Conflict Between Privacy and Personalisation? In *Information Systems Development*. Springer, New York, NY, 1–9. https://doi.org/10.1007/978-1-4614-4951-5_1
- [76] Susanne Trauzettel-Klosinski and Klaus Dietz. 2012. Standardized assessment of reading performance: the new International Reading Speed Texts IReST. *Investigative ophthalmology & visual science* 53, 9 (2012), 5452–5461.
- [77] Horst Treiblmaier and Irene Pollach. 2007. Users' perceptions of benefits and costs of personalization. *ICIS 2007 Proceedings* (2007), 141. <http://aisel.aisnet.org/cgi/viewcontent.cgi?article=1298&context=icis2007>
- [78] TripAdvisor. 2019. Retrieved from <https://www.tripadvisor.co.uk/>. (2019).
- [79] Vasilis Vlachokyriakos, Rob Comber, Karim Ladha, Nick Taylor, Paul Dunphy, Patrick McCorry, and Patrick Olivier. 2014. PosterVote: Expanding the Action Repertoire for Local Political Activism. In *Proceedings of the 2014 Conference on Designing Interactive Systems (DIS '14)*. ACM, New York, NY, USA, 795–804. <https://doi.org/10.1145/2598510.2598523>
- [80] Jayne Wallace, John McCarthy, Peter C. Wright, and Patrick Olivier. 2013. Making Design Probes Work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 3441–3450. <https://doi.org/10.1145/2470654.2466473>
- [81] Yuchen Zhao, Juan Ye, and Tristan Henderson. 2016. The Effect of Privacy Concerns on Privacy Recommenders. In *Proceedings of the 21st International Conference on Intelligent User Interfaces (IUI '16)*. ACM, New York, NY, USA, 218–227. <https://doi.org/10.1145/2856767.2856771>