What's at Stake: Characterizing Risk Perceptions of Emerging Technologies

Michael Skirpan
University of Colorado
Boulder, CO
michael.skirpan@colorado.edu

Tom Yeh
University of Colorado
Boulder, CO
tom.yeh@colorado.edu

Casey Fielser
University of Colorado
Boulder, CO
casey.fiesler@colorado.edu

ABSTRACT

One contributing factor to how people choose to use technology is their perceptions of associated risk. In order to explore this influence, we adapted a survey instrument from risk perception literature to assess mental models of users and technologists around risks of emerging, data-driven technologies (e.g., identity theft, personalized filter bubbles). We surveyed 175 individuals for comparative and individual assessments of risk, including characterizations using psychological factors. We report our observations around group differences (e.g., expert versus non-expert) in how people assess risk, and what factors may structure their conceptions of technological harm. Our findings suggest that technologists see these risks as posing a bigger threat to society than do non-experts. Moreover, across groups, participants did not see technological risks as voluntarily assumed. Differences in how people characterize risk have implications for the future of design, decision-making, and public communications, which we discuss through a lens we call risk-sensitive design.

ACM Classification Keywords

H.1.2 User/Machine Systems: Human Factors; H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords

Risk; Ethics; Big Data; User Attitudes; Mental Models; Design; Emerging Technologies; Big Data; Technology Harm; Algorithms; Filter Bubble; Privacy

INTRODUCTION

Emerging data technologies are primarily developed by capturing or acquiring a large data set to use for further analysis or model training. Designers also build responsive and personalized systems that learn from user behaviors. In both cases, these data are generated from user-specific behaviors tracked and archived, often on public, web platforms such as Facebook, Twitter, etc. Thus, users largely lay the foundations for the accelerating area of Big Data Technologies, including machine learning (ML), artificial intelligence (AI), and behavior-driven design. These users must rely on the companies and parties to whom they have given their data (knowingly or not) to be ethical.

Yet, we already know that many impacts (e.g., privacy, ethical, legal) and constraints (e.g., protocols, technological capabilities) of online technologies are poorly understood by users [24, 8, 36, 15]. We also know that, when asked, users are often uncomfortable or find undesirable the practices of online behavioral advertising (OBA) and personalization [37, 34]. This misalignment is often framed as a consumer trade-off between privacy and personal benefit [13, 40]. Framing it this way leads to an assumption that the benefit of web services must outweigh consumer's privacy concerns since users are not opting out of services.

However, if consumers really are performing this cost-benefit analysis and making a conscious decision, then why we do we see such hype and panic around risks and harms caused by technology in the media? Daily news headlines relay injustice [19, 1, 4, 33], personal boundary violations [32], and gloom [26, 18, 14] over the impacts of technology on society. Some of these problems may indeed warrant concern from the public and social advocates; others might be overblown headlines to keep technology followers clicking. Meanwhile, previous research shows that, when prompted, users do have concerns over issues like data privacy [9, 22, 6]. Therefore, we wonder, do users actually understand the relationship between the data they hand over and the systems being built from that data, including those that could lead to the kinds of risks and harms they read about in the media? What dimensions of how people perceive risk impacts their overall judgment? Are these issues being communicated equally across different pockets of the public?

These questions are also pertinent to understanding the role experts should play in the communication and assessment of risk. Much like a doctor who needs to tell a patient the risks of choosing to do a surgery or not, technology providers and software creators must consider how they communicate possible risks to their users. This role is much different than that of a doctor, given the speed of user consent, the lacking human relationships between engineers and users, and the fact that often services being adopted are free. In this context, we need to better understand how experts and the lay public each think about risk, in order to develop improved expectations around the state of people's internal assessment of risk (i.e.,

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

CHI 2018, April 21–26, 2018, Montreal, QC, Canada

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ISBN 978-1-4503-5620-6/18/04…15.00

DOI: https://doi.org/10.1145/3173574.3173644

their risk perceptions) and then what might be done for better communication, consent, and decision-making.

In order to better understand the psychological dimensions of this problem, we modified and implemented a classic survey instrument from risk perception literature [30]. The original study validated an interesting phenomenon: that the more voluntary one perceives a risk to be, the more willing they are to accept it. The paper also examined the differences between expert and non-expert risk perceptions. We designed a derivative instrument focusing on risks posed to society by emerging technologies (rather than the original study's focus on environmental and health hazards), with a particular emphasis on hazards from technologies emerging due to the rapid availability of data.

After surveying 175 people, 26 of which were technical "experts," we found significant differences in their perception of risks related to technology. We considered an expert someone who was either working in a technical role or at an advanced phase toward earning a degree in a computing field. Our findings indicated interesting contrasts between how the groups ranked and rated risks. We also found that, despite terms of service agreements, many risks related to data technologies were perceived as involuntary by both groups.

These results have implications for design, public communications about technology, policy, and the study of user mental models. In this paper, we discuss the relevance of studying risk perception to HCI, report findings from the deployment of our derivative instrument, and analyze these findings in light of a simple model we are calling "risk-sensitive design" for interpreting when misaligned risk perceptions may warrant reconsideration.

RELATED WORK

Past Work on User Attitudes and Beliefs

A lot of complexity lies in the details of how data is processed and shared once a user opts into a service that captures behavior and information. And while there are benefits for personalized features, numerous studies have shown that users do not always want or like these perceived benefits, or that they do not think these benefits outweigh privacy or related concerns [5, 34]. For example, users think that data being sold to third parties does not benefit them [9].

While some may argue that this is simply a logical trade-off of privacy for convenience or utility [40], research has shown that feelings, comprehension, and reasoning patterns users apply in real situations do not actually support this baseline privacy trade-off. Studies by Cranor et al. have shown that users carry nuanced views about the context in which they feel comfortable with information disclosure, and that they may dislike true data practices, if properly understood [11]. Besmer and Lipford have furthered findings that once misconceptions are clarified, users may regret their technology and disclosure choices [8]. Other studies have shown that it often takes a personal experience before people elevate their awareness of risks they may be ignoring with online technologies [22]. Problems like this could be related to design choices. For example, when given graphics on their mobile devices about when data is collected and where it goes, users are often shocked [6]. Findings by Angulo and Ortlieb point toward the need for better design to educate and inform users around the scenarios they find most concerning [3].

The issue is, we do not always know where these nuanced anxieties truly come from, particularly since, rather than understanding how technologies work, individual users are often applying their own folk models to reason about technology [39, 41]. This makes it difficult to discern ideal forms of transparency in design. We do, however, know that this transparency is not being achieved by the legal agreements that should clarify risk. A number of researchers have shown, for example, that privacy policies are too complex for the average adult [23, 21], and that users often misunderstand what rights they give over to platforms in their content under the copyright policies to which they are bound [17]. Barocas and Nissenbaum further explain [7], that given the complex nature of our data economy, it would be nearly impossible for these fixed agreements to be truly transparent. The legal and technical sophistication of terms of service and data markets make the idea of solving the problem with a unilateral, transparent opt-in agreement a hard goal to chase. However, given the heightened importance of technology and data to all of our lives [35], it is imperative we continue to innovate on how to make users more aware and shape the future of technology to equally benefit everyone.

For these reasons, we believe there is interesting work to be done looking not at what users comprehend from a terms of service, privacy policy, or a particular interface, but instead what risks they perceive related to technology broadly. As research by Acquisti and Gross suggest, privacy concerns may not correspond well to user behavior [2]. Thus, it may be that studying other factors such as risk could clarify dimensions of user reasoning and behavior patterns.

Risk Perception Studies

There have been many interesting studies within the risk perception literature that could be relevant for HCI and technology policy research. The early goals of the field resonate with some the current work studying user attitudes and believes. A foundational paper, "Characterizing Perceived Risk," states the problem as follows: "[This risk perception study] aims to discover what people mean when they say that something is risky; to develop a psychological taxonomy of risk that can be used to understand people's perceptions and predict societal response; and to develop methods for assessing public opinion about risk in a way that is useful for informing policy decisions" [31].

This literature has often tried to characterize how a balance is reached between perceived risks and perceived benefits [28]. We believe this framing may help designers and policymakers working with socio-technical systems. Particularly relevant to HCI, risk perception research shows a higher willingness for people to take on risks when they believe they are voluntary [30]. That is, risks such as driving or football are taken with less concern because the individual has chosen to do so. They also found that the differences in risk perception

between experts and the public lead to concerning trends in inter-group trust and democracy [29].

Seeing these past findings as having relevance to today's conversations about user perceptions of computing systems, we believed there is value in replicating a risk perception instrument that focuses on technologies germane to contemporary society. Next we discuss the shape our adapted instrument took and our methods for seeking and grouping participants.

METHODS

Our interest in applying work from the risk literature to HCI was to garner fresh insights about how users and experts think about the risks posed by new technologies. The original study we replicated came from a 1980 paper called "Facts and Fears: Understanding Perceived Risk" [30]. In the write-up they summarize several studies that worked with a particular instrument for comparatively ranking, individually scoring, and psychological characterizing a long list of risks. With over 1000 citations and a number of validations of prior research in that literature, we were interested to apply it to technologies that are now of high relevance to computing researchers.

The first iteration of their instrument contained 30 risks that participants comparatively ranked from most to least risky, and then also provided a raw risk score on a scale from 1-10 (from "very risky" to "not risky"). Participants made estimates on annual fatalities from the different risks and characterized each risk using 9 psychological factors. A later version of this instrument deepened this analysis by using 90 risks and 18 psychological factors.

Risks included were a mixture of common and technical risks–e.g., smoking, handguns, x-rays, contraceptives, vaccinations, and pesticides. The psychological factors, drawn from prior risk perceptions studies, included voluntariness of the risk, immediacy of effect, knowledge about the risk both individual and by expert scientists, control over the risk, severity of the consequence, the commonality of the risk, and familiarity with the risk. The authors surveyed a group of 15 "experts" who worked in relevant fields such as medicine, risk analysis, and policy, and compared their perceptions to other groups–a sample of students, a sample from the League of Women Voters, and sample from an active outdoors club.

Our adaptation of this survey truncated the size of the instrument–using 18 risks and 6 psychological factors–to keep the survey time to around 20 minutes and reduce the cognitive load required of our participants. We kept in 3 of the original risks for benchmarking and comparison,and then adapted the other 15 risks based on emerging data-driven technologies that have been discussed extensively in research and current events. We used the same psychological factors as the original first iteration of the instrument.

Starting with a list of 30 risks, we shared our initial choices with colleagues working in large tech companies and research institutes, to make sure we chose the most relevant and interesting risks. In the end, the 15 new risks we used were highly relevant to today's reporting on technology issues. "Death or destruction from autonomous drones," "biased algorithms for filtering job candidates," "identity theft", "security breach of an online account," "filter bubbles (individuals receiving different versions of the internet)", "technology divide (technology only benefiting a small elite)", "job loss from automation," "nude photos being leaked," "having your online activities researched without consent," "Distributed Denial of Service Attacks," "Undisclosed third-parties having access to your data," "Online bullying and harrassment," discriminatory algorithms used for policing," "hacktivists leaking large data sets containing personal information," and "malfunctions from driverless cars" were the added risks. From the original survey we kept "nuclear reactor meltdown," "harm to one's health from vaccines," and "plane crashes" which were picked from the top (nuclear reactors), middle (plane crashes), and bottom (vaccines) of average rankings from the original study.

Similar to the original study, participants ranked the 18 risks comparatively and individually. They also characterized the 15 new risks using 6 psychological factors along a 1-7 scale. For our study we used, "voluntariness," "fear of risk," "severity of consequence," "perceived self understanding of risk," "perceived understanding of risk by domain experts," and "likelihood of risk happening," where a 1 minimized the factor (e.g., "involuntary," "do not fear risk at all") and a 7 maximized the factor (e.g., "completely voluntary," "a deep fear of the risk"). In keeping with the goals of the original study, we framed our survey as involving "risks posed to society," and explicitly asked our participants to rank these risks as a threat to people broadly not merely to themselves.

In order to replicate the initial survey's distinction between "experts" and "non-experts," we distributed the survey to multiple populations and made this distinction based on career. For the current study, "experts" constitute people with careers or focuses in computing and technology development. We piloted the survey with both experts and non-experts from our social networks, and iterated on the description of the risks, as well as the wording of survey questions, based on their feedback. We also used pilot responses to tweak the timing of the survey, to ensure that it would take less than half an hour to complete.

Targeting experts, we deployed the survey to people in tech careers via snowball sampling following initial seeding within the first author's social network. To target a general population of non-experts, we also deployed the survey using Amazon Mechanical Turk (mturk), a crowdwork system where workers complete small tasks for payment. Though there are known biases in the mturk population that prevents broad generalizability, it has been shown to have advantages over localized populations for demographic diversity [10]. However, the population does tend to skew younger and less ethnically diverse [20]. Though prior work has shown that mturk workers tend to be slightly more tech-savvy than the general population [16], they also skew less educated overall than most American working adults [20]. In the next section, we report the demographics from both expert and non-expert populations, so that our results can be interpreted with these factors in mind.

We deployed the survey on Mturk in June 2017. Based on our pilot testing, we were aware that the survey took between 15 and 30 minutes to complete. We paid workers $3.00 for the task, ensuring that the rate of pay would typically be over minimum wage. We included two attention checks in the survey (asking for a certain answer, to ensure the questions are being read), a strategy which is used to help ensure the validity of survey responses on mturk [12]. 166 workers completed the survey, and we removed 17 responses for failing attention checks or providing incomplete data, resulting in a total of 175 responses. As part of the survey we also had data about participants' careers. For those who fit our definition of "expert," we removed them from this dataset and added them to the expert dataset. The final result is 26 expert responses and 149 non-expert responses.

In addition to quantitative measures that were drawn from the original risk perception survey instrument we also added two open-response questions to the survey. First, for the three items that they ranked as the riskiest (i.e., the top three), we asked them to describe what they thought the worst case scenario was. Additionally, we asked if there were any additional serious risks to society caused by technology that we had not asked them about. During our analysis stage, we analyzed these responses qualitatively, using iterative, open coding [27].

RESULTS

Analyzing survey responses, our goal was to better understand how people perceive different risks. That is, we view risk as an operational construct within human judgment to reason about certain decisions, such as when to act or not and when a situation is threatening or not. People can therefore only weight these decisions based on what risks they perceive. Our results section starts with an overview our population sample, then highlights group differences for comparative risk ranking, raw risk scoring, and psychological factors as they correlate to perceived riskiness.

Properties of Population Sample

Our sample consisted of 175 completed surveys after removal of participants who (1) did not pass our attention checks, and/or (2) did not complete the survey. Following findings from the original study, we primarily analyzed the difference between expert and non-expert populations. Expert was defined as someone with either a career in computing with a primary role that is technical in nature or a student seeking a degree in a technology-focused field such as computer science or electrical engineering. Using this distinction, we separated our sample into 26 experts and 149 non-experts. The goal of this grouping is to develop an understanding of whether experts perceive technological risks differently than non-experts, given the often complex and future-oriented nature of these risks. Demographically, the sample included 96 men, 77 women, and 2 non-gender identifying individuals. Along ethnic lines, we split our population into 124 subjects who primarily identified as "White" and 51 subjects who primarily identified as a non-white ethnicity. Within groups, demographics were less spread. Our expert group of 26 had 25 men and 1 non-gender-identifying individual. 20 out of the 26 were white with 2 black, 1 hispanic, and 3 mixed-heritage respondents. Our non-expert group contained a wider spread with a nearly 50/50 split between men and women (82 women and 80 men). As expected, our differed widely on educational background. The expert group contained 30% with a masters degree, 46% with a bachelor's degree, 11% with a doctoral degree, and the remaining 3 reported some college (2) or a high school degree (1) but still claiming to work in a technical role. Our non-expert group contained 5% with a masters or professional degree, 44% who had a bachelor's degree, 9% an associates degree, and the remaining 42% having a high school diploma or some college, but no degree.

Though we sought to replicate a prior study that had a small expert group (N=15) to compare against, we were concerned about the cohesion of our expert group given its size (N=26). Though this could be even further improved in future work, we did see somewhat tighter agreement among experts than non-experts. Within items asking respondents to compare risks, non-experts had an average standard deviation of 4.79, while experts had 4.16. On items that asked for a raw 1-10 ranking of risks non-experts had an average standard deviation of 2.76 compared to experts' 2.39. We further noticed a general trend that those in the expert population tended to rank technology risks much higher compared to common risks (e.g., plane crashes, nuclear reactor failure) both comparatively and using a raw score, whereas non-experts were consistently more concerned with these common risks. We will characterize this finding further below, but we took these as signs that there was a coherent enough difference between populations to garner insights.

Risk Ranking and Scoring

As Figure 1 indicates, the comparative ordering of risks (from most to least) was quite different between experts and non-experts. The top of the list looks similar, though non-experts were more concerned with identity theft (p=.063; two-tail t-test) and experts with job loss (p=.068). Average expert rankings showed the top three risks to be (1) job loss, (2) account breach, and (3) identity theft. Non-experts had the same top three items, but in different order: (1) identity theft, (2) account breach, and (3) job loss. This makes sense, given that experts could believe themselves to have a more sophisticated ability to control the identity theft, but not mass job loss. However, it is also notable that these three risks are well documented in media and have a national attention surrounding them.

With respect to participants' 1 to 10 risk score for each item (see Figure 2), the results generally validated comparative rankings; however, we did gain more nuance in terms of the difference in magnitude between certain risks that would have been unclear from rankings alone. Figure 2 plots risk scores comparatively with average expert scores on the x-axis and average non-expert scores on the y-axis. This data shows, in general, non-experts perceive all risks as greater in an absolute sense. With 1 indicating most risky and 10 least, non-experts on average scored all risks below a 6. Experts, on the other hand, scored a third of the list (6 items) at 6 or above. This may suggest there is a more worried perception

Rank	Non-Expert Risk	Non-Expert Mean Rank	Expert Risk	Expert Mean Rank
1	Identity Theft	5.000	Job Loss	5.769
2	Account Breach	6.101	Account Breach	6.385
3	Job Loss	7.678	Identity Theft	6.577
4	Hacktivist Leak	7.980	Technology Divide	6.923
5	Auto-Drones	8.523	Bias Job Alg	7.192
6	Harassment	9.074	Discriminatory Crime Alg	7.231
7	Undisclosed third party	9.349	Hacktivist Leak	7.231
8	DDoS	9.403	Filter Bubble	7.654
9	Nuclear Reactor Meltdown	9.644	DDoS	8.269
10	Discriminatory Crime Alg	9.758	Undisclosed third party	8.462
11	Research w/o Consent	10.141	Harassment	9.346
12	Bias Job Alg	10.154	Auto-Drones	9.808
13	Driverless Car Malfunction	10.315	Research w/o Consent	11.154
14	Technology Divide	10.765	Nude Photos	12.038
15	Plane Crash	11.060	Driverless Car Malfunction	12.269
16	Filter Bubble	11.362	Nuclear Reactor Meltdown	14.308
17	Nude Photos	11.846	Plane Crash	14.654
18	Vaccine	12.846	Vaccine	15.731

Figure 1. Average comparative risk ranking by non-experts vs experts where items with significant differences (p<.05 for two-tailed t-test) are highlighted.

of technology, broadly, in non-expert risk judgments. Though it is also possible that experts do not judge the broader risks of technology with enough gravity. Research without consent, while being only slightly different when comparatively ranked, showed a difference of 1.237 in mean score (p=.015) with non-experts ranking this as riskier than experts. This is not so surprising, though it is salient given the lively discussion currently ongoing around online research ethics [38].

Psychological Factors of Risk

We now move to examining how the psychological characteristics drawn from the original study played into perceived riskiness on the raw 1-to-10 scale. In our analysis, we were interested in the averages of factors that made up the psychological space surrounding a risk together with the individual correlations between a particular factor and its relationship to a risk score. Therefore, positive correlations reported below approximate the amount that a certain psychological factor influences an overall riskier perception of the associated item. Inversely, negative correlations indicate psychological factors that weigh in towards a lower perception of overall risk.

Our study broadly validated prior findings in the risk literature that show the more voluntary a risk is perceived to be, the less risky it is ranked and scored. For all except two items, higher perceived voluntariness correlated negatively with perceived riskiness. This is to be anticipated as, much like our acceptance of cars despite their danger, we expect people to take on more risk when they have chosen this risk and perceive themselves as in control. The two exceptions we saw made some sense – DDoS attacks and Death and Destruction from Autonomous Drones which, only for non-experts, voluntariness had a minor positive correlation with risk scores (r=.05 and .1, respectively). Given that there is really no voluntary dimension to either of these risks, it is unsurprising that voluntariness was basically uncorrelated even if slightly positive. If properly understood, the mild positive relationship may even be from a rightful recognition that owning insecure devices that allow DDoS to happen or not actively pushing to downgrade our military and suppress drone warfare are both voluntary choices that expand these risks.

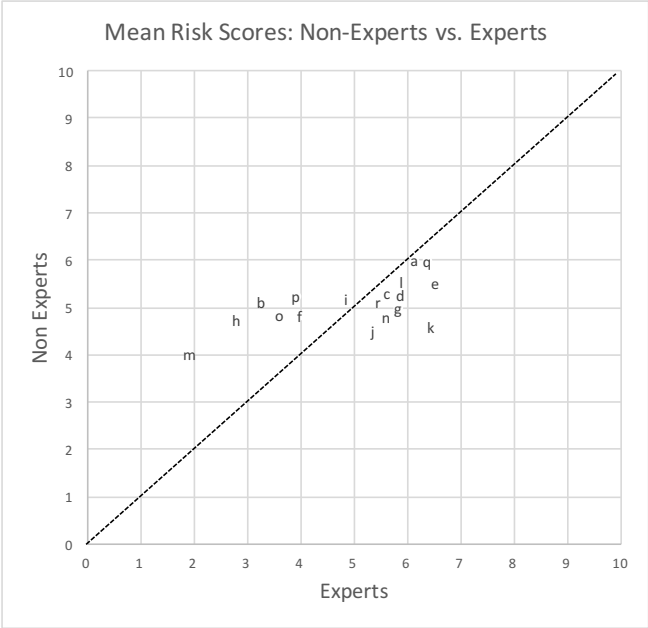

a. Identity Theft	g. DDoS	m. Vaccine
b. Nuclear Meltdown	h. Plane Crash	n. Tech Divide
c. Harassment	i. Auto-Drone	o. Research w/o consent
d. Discriminatory Crime Alg	j. Bias Job	p. Driverless Car Malfunction
e. Job Loss	k. Filter Bubble	q. Account Breach
f. Nude Photos	l. Hacktivist Leak	r. Undisclosed third party

Figure 2. Risk perception by experts vs non-experts for 18 technologies.

However a novel complication from the prior risk literature is that both groups rated nearly all risks related to emerging technologies as characteristically involuntary. That is, using our 1-7 Likert Scale, almost every risk carried an average score below 3 and in many cases below 2. This suggests that despite the consent processes built into most software and web services, the corresponding risks are not perceived as something being voluntarily assumed. There were a few minor exceptions where voluntariness ranked slightly higher. Within the non-expert grouping two risk items ended up with a mean voluntariness rating above 3: driverless car malfunction (3.43) and nude photographs (3.89) which were still below the middle of our scale. The expert list was the same – driverless cars (3.5) and nude photos (4.15) – with the addition of filter bubble (3.73). This is a good validation within our findings since among our items these are all risks that do have some immediate control by individuals – not getting in a driverless car, never letting nude photographs be digitized, and actively seeking news and information from sources outside of your ideological affiliations.

Other psychological factors of importance were perceived fear and severity, which both often carried a highly positive correlation with perceived risk. The perceived riskiness of destruction from autonomous drones, driverless car malfunctions, research without consent, and hacktivist leaks all had their strongest positive correlations with either fear or severity across both groups. This might imply that more than facts or reasoning some risks get a status due to the imagined consequences being frightening or intensely concerning. This

may explain the somewhat surprising result that, even for experts, the perceived riskiness of driverless car malfunctions was most positively correlated with fear (r=.717), which outweighed the strongest negative correlating factor of the perceived understanding of domain experts (r=-.47). This suggests that the innate fear of the autonomous car crashing factors higher in one's risk calculation than their belief that experts are on top of debugging and testing.

Another notable result was that the expert group showed a positive correlation (r=.4) between how well domain experts understand the riskiness of hacktivist leaks–suggesting that the more experts understand the domain, the riskier it becomes. This is in contrast to most of the other technology risks, where more expert knowledge makes something less risky. This finding may imply that as experts get better at understanding data breaches and leaks that they too could take part in hacktivist activities, or that even highly qualified experts cannot stop the ability for insiders and creative hackers to gain access to secrets. Other risks where expert understanding had a positive correlation (above r=.2) with perceived risk were discriminatory crime algorithms (r=.27) and filter bubble (r=.26). These risks probably warrant more niche domain expert opinions since they are both still new and ill understood by technologists broadly. It seems there is some belief that experts being on top of it might lower the risk, but it is not yet a strong one.

Figures 3 shows a graph that compares psychological factors between groups on two items that show significant differences between experts and non-experts: research without consent and filter bubble. As the graphs signify, experts appear to be generally more worried about the filter bubble whereas non-experts are more concerned about research without consent. Interestingly, experts have a fairly strong positive correlation between self understanding and perceived riskiness (r=.53) and similarly with likelihood of occurrence and perceived riskiness (r=.517). This finding could be straightforward, given that experts are probably more likely to have taken time to comprehend and break out of their own filter bubbles due to the current events and recent impacts from the issue. On the other hand, fear (experts r=.48; non-experts r=.295) and severity (experts r=.476; non-experts r=.269) have the highest positive correlations for both groups concerning research without consent. Though we cannot say definitively what exact fears they harbor or severe impacts they imagine, we will discuss in the next section some of the "worst case scenarios" that some participants stated with respect to these risks. Regardless, our findings do suggest that people believe there are real risks to allowing online data to flow into research without consent. Public perceptions of this domain could be heavily influenced by media explanations of recent controversies such as the Facebook emotional contagion study [25] and the public release of OKCupid data by a researcher [43].

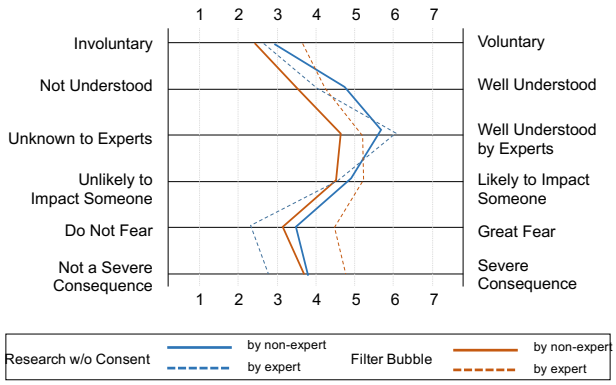

Figure 3. Comparing the psychological factors regarding the filter bubble (orange) vs research without consent (blue) as perceived by non-experts (dotted) vs by experts (solid).

Worst Case Scenarios

In addition to the psychological factors, asking participants about their idea of worst case scenarios for the technologies they consider riskiest gives us a more qualitative sense of what specific kinds of risks they consider to be the worst. Every technological risk was ranked within the top three by at least ten people.

We broadly identified five primary types of harm that we used to categorize the worst case scenario responses. These included: financial, legal, reputational, physical harm or death, and societal (see Table 1 for examples of each).

Based on the aggregate rankings of risk (see Figure 1), the top three technologies perceived as riskiest by both experts and non-experts are seen as financial and legal risks. Worst case scenarios for identity theft include "losing everything and being imprisoned" and "someone wipes out my entire checking account." Job loss from automation evokes fear of "widespread unemployment and the collapse of our economic system" and for account breach, "credit cards being used to make unauthorized purchases." Financial is also the most common type of perceived risk, also dominating responses biased algorithms, collection of public information, and data sold to third parties.

With respect to differences between experts and non-experts, technologies that experts rate as significantly riskier than non-experts (technology divide, biased job algorithms, discriminatory crime algorithms, and filter bubble) all have dominantly "societal" worst-case scenarios. For example, "bringing down critical services would cause mass chaos" (DDOS), "extremely polarized news and information to the point where no one is certain of truth" (filter bubble), and "extreme divides in economic opportunities and outcomes, leading to social collapse" (technology divide). One possible explanation for this finding is that more experts are able to foresee the broader impacts of technology, than non-experts.

Another aspect of risk perception that these free responses highlights is that participants often have very different ideas of concrete risks for these technologies. For example, for DDOS, most participants mentioned worst case scenarios like bringing down emergency services or energy grids, but one wrote, "People being angry that they canâĂŹt get on Facebook." Similarly, a number of participants, mostly among non-experts, thought that the risks around filter bubbles had

Scenario Type	Risk	Example Quote
Financial	DDoS	"bank account drained of money"
Legal	Discriminatory Crime Algorithm	"being imprisoned for nothing"
Reputational	Nude Photos	"massive embarrassment and you lose your job"
Physical Harm or Death	Online Harrassment	"it slowly erodes you, so suicide would be the worst"
Societal	Hacktivist Data Leaks	"civil unrest making us vulnerable"

Table 1. Scenario Types, Risks, and Example Text from Respondents' Open-Ended Worst-Case Scenario Response

to do with government controlled or censored Internet. For example, "They could choose what to let you know about what not to know so the earth could be under alien attack and they could just block that if they wanted to." The risk for online research without consent also showed a number of conflicting ideas of worst case scenarios, particularly among non-experts. Whereas some mentioned "invasion of privacy" and similar, which tracks to existing work around reactions to online research [42], we also saw what appears to be misunderstandings of what research of public data actually entails–for example, "getting lots of spam calls" and "could lead to stalking by the government."

Additional Risks

We also asked participants in an open response question what they saw as any other major technological risks to society. Less than a quarter of our participants included a response, which suggests that our list of risks was fairly comprehensive. Some of these were more specific or nuanced versions of risks we assessed–for example, "leakage of personal data through apps" or "hackers causing voter fraud." Others were technological risks not as related to computers or data–for example, "biotoxins" or "solar flares wiping out technology."

The theme that appeared most frequently among those who answered this question was addiction to or reliance on technology. Example answers were "young people of future generations will be lost without tech," "addition to technology," and "inability to get along without technology." In future work, this would be a good issue for further exploration, since it is clearly a risk that is on people's minds.

DISCUSSION

Tradeoffs and Voluntariness

Without an understanding of risk perception, researchers and technologists could be missing critical information about decision-makers and users alike when making choices about regulation, technological design, or technology adoption. It is equally important to understand the perceptions of experts and users. Critically, we should be prepared for either group to have misperceptions. While we all hope experts are able to objectively assess their own products, it is important to acknowledge the possibility of bias or misunderstanding in their perceptions. Thus, we should not see either expert or user mental models as a gold standard for true risk, but having a higher-level view of each together tell us a lot about the broader scope of risk perceptions and where issues may be laden.

Our findings around voluntariness are particularly salient when it comes to user models of risk. It is a robust finding in risk literature that risks perceived as voluntary (such as driving and skiing) are seen as more acceptable to the public even if statistically they are dangerous [30]. When modeling user and consumer decisions around data-driven and networked technologies, a similar schema is often deployed to explain trading off privacy for convenience. There is an assumption that since users accept the terms of service and then provide information and content, any repercussions are part of the agreement and thus voluntarily assumed, despite a great deal of research having shown that it is unlikely that users read and understand terms of service [23, 21]. Similarly, releasing new features or adopting new algorithms that shape behavioral and social patterns are not considered problematic because use of the tool is voluntary. Our findings complicate this assumption since, in nearly all cases, technological risks are not seen as voluntary. This might illuminate the media addiction to reporting on technology's mishaps, since people may not really see these risks as something for which they signed up.

Given complementary findings around users' lacking understanding of the legal agreements to which they are bound, this perception of involuntariness might be expected despite common assumptions. However, this trend is actually more worrying since our expert participants generally ranked complex technological risks higher. This gap in expert and non-expert perceptions implies that users may not only feel they are being involuntarily harmed by the choices technologists make with data and features development, but that they may not even really realize what is at stake in the coming years. Taking this finding seriously, designers may want to offer more information and choices about how features may be affecting user experience. Technologists broadly could attempt to allow for more discussion, feedback, and disagreement around new features and data practices, and perhaps be willing to hold off rolling out a feature when it raises more concerns than excitement.

Our exceptional cases to this finding actually further validate this take. Given that nude photos being shared without consent is seen as more voluntary than the other risks we examined, we can speculate that this perception is based on interpreting it as a voluntary choice to take or share nude photos at all. And while our results do not indicate it is perceived as fully voluntary, since the implication is the photos go beyond where they were meant to be seen, this does indicate an attitude that the burden is on the user not to digitize and share nude photographs. Similarly, with malfunctions in driverless cars, which may at first glance seem surprising to be viewed as a voluntary risk, we can speculate that this perception is based on seeing a choice to ride in a driverless car.

Most Technologies Are Riskier Than You Think

We also found that experts generally rank technological risks as comparatively more risky than non-experts. One takeaway we suggest from this finding is that it would be in society's best interest for researchers to focus more acutely on the complexity of education and public communications. With algorithms and AI mediating more of our relationships and institutions, yet being highly esoteric topics, it is concerning that the risks of these changes are perceived as stronger by the group who is involved in the creation of those risks. As we have already seen with issues such as climate change, vaccines, and immigration, once complex problems come to the attention of large swaths of the public who do not understand them, it can catalyze division, exploitation, and people acting out of the best interest of our society.

Some of these risks are also important for non-experts to understand as we move into new futures suggested by technology–for example, as AI leads to the automation of more tasks and potential job loss. It is notable that both groups rated job loss in their top 3 risks; yet, it is also interesting that experts rated this #1 versus non-experts #3. If there were to be a mass shift in jobs, e.g., from autonomous trucking and shipping, it is particularly important that this change is well communicated and comes with a policy plan for mitigation. Otherwise, we may be steering toward another crisis emphasizing lack of trust between experts and the public. The potential for such a problem expands when considering that our findings suggest that non-experts are not taking the possibility of bias or discriminatory algorithms as seriously as are experts. Following workshops such as FATML (Fair, Accountable, and Transparent Machine Learning), technologists should take a stance on what makes an algorithm fair and communicate these standards to their users. Designers may also strive to make users aware of when they are being acted on by an algorithm and promote different kinds of reflection when displaying results.

What Makes People Afraid?

Our psychological factor analysis revealed that, besides voluntariness, fear and severity play heavily into people's perception of risk (both expert and non-experts). This is a challenging result to interpret without a deeper depiction of the imagined consequences. We suspect that fear has some interaction with voluntariness, as we often think of fears as having a somewhat irrational component to them as they are personal and can be uncontrollable. Such an explanation might help us understand why the risk of malfunctioning driverless cars was highly relative to the person's fear regardless of expert understanding.

Severity is even harder to analyze as this relates to the specifics of the imagined consequence, especially since some consequences even experts could not fully comprehend yet. We do know from our analysis of free response answers regarding worst case scenarios, that people's imagined consequences can vary considerably. For example, research without consent carried with it differing ideas of what this might mean–some non-experts think that it could relate to receiving spam or to government surveillance, which suggests a misunderstanding of what research on public data entails. Others may have a better idea of potential consequences, imagining that they get pooled and identified in categories to which they do not believe they should belong. Fear could also relate to some deeper belief about what it means to be studied. Regardless, the perceived risk around this practice, whether rational or not, is something that researchers should keep in mind when collecting data.

With respect to experts' perceived fear, another finding that supports our understanding of current events is the significant difference in how experts versus non-experts see the filter bubble. Not only did experts rank the issue as comparatively more risky, but also many of their judgments of psychological factors were quite different. As public opinion and the media often blame fake news and political polarization on technologists, it is important that we get a better sense of how this problem is understood by the public. One might wonder if the filter bubble is truly just perceived as less risky by non-experts or if it is another risk that is poorly understood due to its technical provenance.

Our analysis of worst case scenario responses suggests that it could be the latter, since a number of participants framed the problem as being one of purposeful government control of the Internet. However, regardless of causality, it is also concerning that beliefs about self understanding of the problem correlated highly with experts believing it was risky. This suggests that the underlying problem may be situated in an educational space. As technologists deliver new personalized features to users, we may only amplify the problems surrounding the filter bubble without corresponding communications around how to work around or see outside of personalized experiences. Otherwise, given our results, it may be the case that people do not fully see how separate the lived realities of people across the country and world truly are and thus this problem is not taken seriously. Given other research that suggests, when asked, people are often skeptical of personalization [2], the fact that people do not see filter bubbles as a high risk may suggest a new challenge for designers: how to make transparent when and how personalized filtering is happening, and what the consequences are.

RISK-SENSITIVE DESIGN

Taken together, our results suggest a number of thoughtful approaches to public communications, policy, and technology design. In addition to the implications already discussed, we also propose a design principle that we term—**risk-sensitive design**. Risk-sensitive design recommends that design decisions regarding risk mitigation features for a particular technology should be sensitive to the difference between the public's perceived risk and the "acceptable marginal perceived risk" at that risk level. One must remember that technologies and risks are not always one-to-one and it may be a particular system or design that is creating a negative consequence rather than an individual technology.

We developed a graphical method as a tool to illustrate this design principle. Figure 3 plots risk perception by experts vs non-experts for the 18 types of risks we studied. Figure 4 is an abstract version of the same plot, that we will use to explain

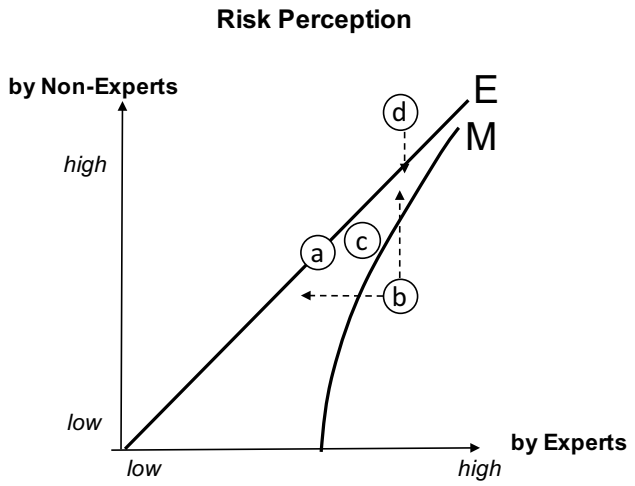

Figure 4. An estimated Risk Perception Curve. Risk-Sensitive Design proposes that risk mitigation strategies should be informed by the difference between public and experts' risk perception and the degree to which such difference is acceptable ($E - M$).

risk-sensitive design. Each circle represents a risk presented by a technology for which its perception level (how risky each population perceives the risk to be) has been assessed. On the diagonal line, experts and non-experts perceive risk the same way. For brevity, we will refer to this line of equal risk perception as E.

For example, risk a is perceived by non-experts as no more and no less risky, compared to experts. Above the diagonal line, non-experts perceive risks higher than do experts, such as technology d. Below the diagonal line, non-experts perceive risks lower than do experts, such as risks b and c. The downward bending curve represents the concept of "acceptable marginal perceived risk." For brevity, we will use M to denote this curve. We introduce this curve in order to push for an important design heuristic—"how much can the public underestimate a risk before it is deemed unacceptable?" The higher the risk, the smaller the margin should be, which is modeled by the downward bend.

Observe that the two curves E and M divide the graph into three regions. Risk-sensitive design argues for different design recommendations based on the region in which a technology lies. We discuss each region in turn.

Most important, in our opinion, is the region represented by risk b. Here, the non-expert public grossly underestimates the risk b associated to some technology. This is reflected by its position underneath the acceptable marginal perceived risk curve (M). As shown, the gap between M and E is large and may be instigate broader harms than can be mitigated by simply refining these features. We take a provocative stance to recommend the underlying technology or technical arrangement should be **reconsidered immediately**. Meaning, studying the impacts the technology is having, communicating to users the concerns related to the technology, working in consortium with others to elevate awareness and adopt policies toward mitigation, and potentially even recalling or scaling back the technology until the risk is better understood. Right now, given our findings, personalization technologies applied to content feeds relevant to public decision-making, which in turn cause filter bubbles, may lie in this region of concern.

We believe keeping a technology found to be creating risks in this region out there with no change or revision could be doing more harm than good to the society. While a technology is under review or even temporarily recalled, there are two potential actions. First, scientists and engineers can further innovate to reduce the risk exposure b caused by the technology, which would then be reflected in the graph by a leftward movement. Second, more investment can be made to increase public awareness of risk b, which is reflected in the graph by an upward movement. Both offer a path to exit this questionable region and enter the relatively safer region where technology c is. Once in a safer region, the associated technology can be re-introduced in its new form. We discuss this relatively safer region next.

The region flanked by M and E is considered relatively safer, but not entirely safe. Technology c's risk is still underestimated by the non-expert public. However, the gap is small enough to be within an acceptable range (above M). We recommend **strong safety and risk mitigation measures**. Depending on the kind of risk, examples might include two-factor authentication, compulsory user safety training, or changes to business practices or design. Additionally, better communication or education could also help reduce this gap. In some cases, tactics such as showing real-world examples of harm could be desirable.

The region above E is where a technology, such as d, is overly perceived as risky by the non-expert public. There is less concern that a person could be harmed as a result of not being cautious enough or misinformed. However, this situation could lead to diminished use of the technology or over-reactions when complications do occur. We recommend **adopting a communication strategy focusing on reducing fear and misconceptions**. The aim is to reduce the perception of risk d, as represented by the downward arrow.

This paper contributes knowledge of where on the plot the 18 types of (mostly technical) risk lie. This paper further argues that a line ought be drawn for the acceptable marginal perceived risk (M) and offers design recommendations based on this line. But as to *where* this line should be drawn, we believe it is a subject of further public discourse, which we hope this study and model provoke.

Decision-Making with Risk-Sensitive Design

At the time we formulated the principle of risk-sensitive design, we had not yet completed our survey and did not know which technology risks would fall into the lower region where we would recommend rethinking the technology. Through our analysis, the two risks that fell closest to the lower-right region are filter bubbles caused by personalized recommendation systems and job loss caused by automation and AI. Both of these are pervasive and intertwined with many other benefits industry and research are likely less inclined to put on hold, or that are challenging to rethink. Deciding where

to draw these lines is a difficult task, and proposing that certain technologies might be unacceptable within the bounds of risk-sensitive design is likely to generate debate. Unacceptable, in this context should be seen as a recognized red flag where the vast difference between engineers understanding in the public may lead to misuse or harm. The point being to highlight where we might pause and do deeper study.

Thus, we feel that more important than using this tool to make concrete judgments about a particular technology, is using it to provoke the right set of questions. For example, our qualitative findings show that filter bubble is a technology that is poorly understood by our participants. This suggests that, though experts rank it as risky, the gap might be more easily closed not by attempting to make the technology less risky but by educating the public about how it works and what the risks are. With respect to job loss from automation, this risk is co-evolving with the continued advancement of technology, and should be approached with a combination of risk mitigation in design and public communication and education to reduce the expert/non-expert gap. Beyond design, this finding implicates the need for policy considerations.

Risk-sensitive design calls for designers of new technologies to go beyond traditional risk and benefit analysis and to also pay attention to how their users may perceive risk, provoking thoughtfulness about how to introduce a technology in a socially responsible way. One should bear in mind the question of "where might risks created by my technology enter into this space" and make efforts not to prematurely introduce a technology if evidence projects an entry point below *M* in our graphical tool.

We also recommend that risk perception should be included as a factor to design as early in the design process as possible, rather than as an afterthought once a technology is already built and deployed. For example, focus groups and field observations in formative studies should include risk perception as a factor, not just an objective risk assessment. Summative evaluation should also include instruments to measure risk perceptions with study participants, and multiple expert opinions should be weighed as well. The survey instrument we used is one model for obtaining risk perception data.

One limitation of this approach is that we assume risk is known and can be predicted by experts, as a measurement alongside user perceptions. For some technologies, however, risk is not known until it is introduced. Future expansion of our model could include uncertainties, such as modeling a technology's position in the graph as a probabilistic bubble rather than a dot, and movement within the space as a funnel rather than a line.

CONCLUSION

The above study applied an instrument from the risk perception literature to analyze thinking about emerging technologies. We found that generally users do not think of risk exposure from technology to be voluntary. There were also considerable differences in how risk is perceived by experts and the lay public, which may explain why problems such as the filter bubble, have become so concerning. Our paper ends with a discussion of risk sensitive design, hoping to provoke continued conversation about how technologists should respond when there are large gaps in how the public and experts think about risk. We hope to see HCI researchers continue study in this area and advance the conversation further.

REFERENCES

1. 2015. Google apologises for Photos app's racist blunder. *BBC News* (July 2015). http://www.bbc.com/news/technology-33347866
2. Alessandro Acquisti and Ralph Gross. 2006. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *International workshop on privacy enhancing technologies*. Springer, 36–58.
3. Julio Angulo and Martin Ortlieb. 2015. "WTH..!âĄĹ" Experiences, reactions, and expectations related to online privacy panic situations. In *Symposium on Usable Privacy and Security (SOUPS)*. https://www.usenix.org/system/files/conference/soups2015/soups15-paper-angulo.pdf
4. Julia Angwin, Lauren Kirchner, Surya Mattu, and Jeff Larson. 2016. Machine Bias: There's Software Used Across the Country to Predict Future Criminals. And it's Biased Against Blacks. (May 2016). https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing
5. Annie I. Anton, Julia B. Earp, and Jessica D. Young. 2010. How internet users' privacy concerns have evolved since 2002. *IEEE Security & Privacy* 8, 1 (2010). http://ieeexplore.ieee.org/abstract/document/5403147/
6. Rebecca Balebako, Jaeyeon Jung, Wei Lu, Lorrie Faith Cranor, and Carolyn Nguyen. 2013. "Little Brothers Watching You": Raising Awareness of Data Leaks on Smartphones. In *Proceedings of the Ninth Symposium on Usable Privacy and Security (SOUPS '13)*. ACM, New York, NY, USA, 12:1–12:11. DOI: http://dx.doi.org/10.1145/2501604.2501616
7. Solon Barocas and Helen Nissenbaum. 2014. Big data's end run around anonymity and consent. *Privacy, big data, and the public good: Frameworks for Engagement* (2014), 44–75.
8. Andrew Besmer and Heather Richter Lipford. 2010. Users'(mis) conceptions of social applications. In *Proceedings of Graphics Interface 2010*. Canadian Information Processing Society, 63–70. http://dl.acm.org/citation.cfm?id=1839226
9. Igor Bilogrevic and Martin Ortlieb. 2016. "If You Put All The Pieces Together...": Attitudes Towards Data Combination and Sharing Across Services and Companies. ACM Press, 5215–5227. DOI: http://dx.doi.org/10.1145/2858036.2858432
10. Buhrmester, Michael, Kwang, Tracy, and Gosling, Samuel D. 2011. Amazon's Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data? *Perspectives on Psychological Science* 6, 1 (Jan. 2011), 3–5. DOI:http://dx.doi.org/10.1177/1745691610393980

11. Cranor, Lorrie F., Reagle, Joseph, and Ackerman, Mark S. . 2000. Beyond Concern: Understanding Net Users' Attitudes about Online Privacy. In *The Internet Upheaval: Raising Questions, Seeking Answers in Communications Policy*, Ingo Vogelsang and Benjamin M. Compaine (Eds.). MIT Press, Cambridge, Massachusetts, 47–70.

12. David J. Hauser and Norbert Schwarz. 2016. Attentive Turkers: MTurk participants perform better on online attention checks than do subject pool participants. *Behavior Research Methods* 48, 1 (March 2016), 400–407. DOI: http://dx.doi.org/10.3758/s13428-015-0578-z

13. Serge Egelman. 2013. My Profile is My Password, Verify Me!: The Privacy/Convenience Tradeoff of Facebook Connect. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2369–2378. DOI:http://dx.doi.org/10.1145/2470654.2481328

14. Mostafa El-Bermawy. 2016. Your Filter Bubble is Destroying Democracy. *WIRED* (Nov. 2016). https://www.wired.com/2016/11/filter-bubble-destroying-democracy/

15. Mauricio S. Featherman and John D. Wells. 2010. The Intangibility of e-Services: Effects on Perceived Risk and Acceptance. *SIGMIS Database* 41, 2 (May 2010), 110–131. DOI: http://dx.doi.org/10.1145/1795377.1795384

16. Casey Fiesler, Michaelanne Dye, Jessica L. Feuston, Chaya Hiruncharoenvate, C.J. Hutto, Shannon Morrison, Parisa Khanipour Roshan, Umashanthi Pavalanathan, Amy S. Bruckman, Munmun De Choudhury, and Eric Gilbert. 2017. What (or Who) Is Public?: Privacy Settings and Social Media Content Sharing. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. ACM, New York, NY, USA, 567–580. DOI:http://dx.doi.org/10.1145/2998181.2998223

17. Casey Fiesler, Cliff Lampe, and Amy S Bruckman. 2016. Reality and Perception of Copyright Terms of Service for Online Content Creation. ACM Press, 1448–1459. DOI: http://dx.doi.org/10.1145/2818048.2819931

18. Samuel Gibbs. 2015a. Musk, Wozniak and Hawking urge ban on warfare AI and autonomous weapons. *The Guardian* (July 2015). https://www.theguardian.com/technology/2015/jul/27/musk-wozniak-hawking-ban-ai-autonomous-weapons

19. Samuel Gibbs. 2015b. Women less likely to be shown ads for high-paid jobs on Google, study shows. *The Guardian* (July 2015). https://www.theguardian.com/technology/2015/jul/08/women-less-likely-ads-high-paid-jobs-google-study

20. Paul Hitlin. 2016. Research in the Crowdsourcing Age, a Case Study. (July 2016). http://www.pewinternet.org/2016/07/11/research-in-the-crowdsourcing-age-a-case-study/

21. Carlos Jensen and Colin Potts. 2004. Privacy Policies As Decision-making Tools: An Evaluation of Online Privacy Notices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 471–478. DOI: http://dx.doi.org/10.1145/985692.985752

22. Jennifer King, Airi Lampinen, and Alex Smolen. 2011. Privacy: Is there an app for that?. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*. ACM, 12. http://dl.acm.org/citation.cfm?id=2078843

23. 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. DOI:http://dx.doi.org/10.1145/2470654.2481371

24. Aleecia McDonald and Lorrie Faith Cranor. 2010. *Beliefs and Behaviors: Internet Users' Understanding of Behavioral Advertising*. SSRN Scholarly Paper ID 1989092. Social Science Research Network, Rochester, NY. http://papers.ssrn.com/abstract=1989092

25. Gregory S. McNeal. 2014. Facebook Manipulated User News Feeds To Create Emotional Responses. (June 2014). https://www.forbes.com/sites/gregorymcneal/2014/06/28/facebook-manipulated-user-news-feeds-to-create-emotional-contagio

26. Claire Cain Miller. 2017. Evidence That Robots Are Winning the Race for American Jobs. *The New York Times* (March 2017). https://www.nytimes.com/2017/03/28/upshot/evidence-that-robots-are-winning-the-race-for-american-jobs.html

27. Johnny Saldana. 2015. *The coding manual for qualitative researchers*. Sage.

28. Paul Slovic. 1987. Perception of Risk. *Science* 236 (April 1987), 280–285. http://www.heatherlench.com/wp-content/uploads/2008/07/slovic.pdf

29. Paul Slovic. 1993. Perceived risk, trust, and democracy. *Risk analysis* 13, 6 (1993), 675–682.

30. Paul Slovic, Baruch Fischhoff, and Sarah Lichtenstein. 1980. Facts and fears: Understanding perceived risk. *Societal risk assessment: How safe is safe enough* 4 (1980), 181–214.

31. Paul Slovic, Baruch Fischhoff, and Sarah Lichtenstein. 1985. *Characterizing Perceived Risk*. SSRN Scholarly Paper ID 2185557. Social Science Research Network, Rochester, NY. https://papers.ssrn.com/abstract=2185557

32. Olivia Solon. 2017. 'This oversteps a boundary': teenagers perturbed by Facebook surveillance. *The Guardian* (May 2017). https://www.theguardian.com/technology/2017/may/02/facebook-surveillance-tech-ethics

33. Latanya Sweeney. 2013. Discrimination in online ad delivery. *Queue* 11.3 (2013), 10. http://dataprivacylab.org/projects/onlineads/1071-1.pdf

34. Eran Toch, Yang Wang, and Lorrie Faith Cranor. 2012. Personalization and privacy: a survey of privacy risks and remedies in personalization-based systems. *User Modeling and User-Adapted Interaction* 22, 1-2 (April 2012), 203–220. DOI: http://dx.doi.org/10.1007/s11257-011-9110-z

35. Zeynep Tufekci. 2015. Algorithmic Harms beyond Facebook and Google: Emergent Challenges of Computational Agency. *Colorado Technology Law Journal* 13.2 (2015), 203–218. http://ctlj.colorado.edu/wp-content/uploads/2015/08/Tufekci-final.pdf

36. Joseph Turow, Lauren Feldman, and Kimberly Meltzer. 2005. Open to Exploitation: America's Shoppers Online and Offline. *A Report from the Annenberg Public Policy Center of the University of Pennsylvania* (2005). https://works.bepress.com/joseph_turow/10/

37. Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising. In *Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS '12)*. ACM, New York, NY, USA, 4:1–4:15. DOI: http://dx.doi.org/10.1145/2335356.2335362

38. Jessica Vitak, Katie Shilton, and Zahra Ashktorab. 2016. Beyond the Belmont Principles: Ethical Challenges, Practices, and Beliefs in the Online Data Research Community. ACM Press, 939–951. DOI: http://dx.doi.org/10.1145/2818048.2820078

39. Rick Wash. 2010. Folk models of home computer security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*. ACM, 11. http://dl.acm.org/citation.cfm?id=1837125

40. Mu Yang, Yijun Yu, Arosha K. Bandara, and Bashar Nuseibeh. 2014. Adaptive sharing for online social networks: a trade-off between privacy risk and social benefit. In *Trust, Security and Privacy in Computing and Communications (TrustCom), 2014 IEEE 13th International Conference on*. IEEE, 45–52. http://ieeexplore.ieee.org/abstract/document/7011232/

41. Yaxing Yao, Davide Lo Re, and Yang Wang. Folk Models of Online Behavioral Advertising. (????).

42. Michael Zimmer. 2010. "But the data is already public": on the ethics of research in Facebook. *Ethics and Information Technology* 12, 4 (Dec. 2010), 313–325. DOI:http://dx.doi.org/10.1007/s10676-010-9227-5

43. Michael Zimmer. 2016. OkCupid Study Reveals the Perils of Big-Data Science. (May 2016). https://www.wired.com/2016/05/okcupid-study-reveals-perils-big-data-science/