Privacy Norms and Preferences for Photos Posted Online

ROBERTO HOYLE, Oberlin College

LUKE STARK, Microsoft Research

QATRUNNADA ISMAIL, King Saud University

DAVID CRANDALL and APU KAPADIA, Indiana University Bloomington

DENISE ANTHONY, University of Michigan

We are surrounded by digital images of personal lives posted online. Changes in information and communications technology have enabled widespread sharing of personal photos, increasing access to aspects of private life previously less observable. Most studies of privacy online explore differences in individual privacy preferences. Here we examine privacy perceptions of online photos considering both social norms, collectively—shared expectations of privacy and individual preferences. We conducted an online factorial vignette study on Amazon’s Mechanical Turk ($n = 279$). Our findings show that people share common expectations about the privacy of online images, and these privacy norms are socially contingent and multidimensional. Use of digital technologies to share personal photos is influenced by social context as well as individual preferences, while such sharing can affect the social meaning of privacy.

CCS Concepts: • **Security and privacy** → **Social aspects of security and privacy**; • **Human-centered computing** → **Social content sharing**;

Additional Key Words and Phrases: Privacy, image sharing, contextual integrity

ACM Reference format:

Roberto Hoyle, Luke Stark, Qatrunnada Ismail, David Crandall, Apu Kapadia, and Denise Anthony. 2020. Privacy Norms and Preferences for Photos Posted Online. *ACM Trans. Comput.-Hum. Interact.* 27, 4, Article 30 (August 2020), 27 pages.

<https://doi.org/10.1145/3380960>

1 INTRODUCTION

Today we are surrounded by digital images. Unlike the pervasive presence of advertising and broadcast media images in the 20th century, today’s ubiquitous images are part of our personal lives, including everything from friends at a party, babies, and foreign travel, to more novel types

This material is based upon work supported by the National Science Foundation under grants CNS-1408730 and CNS-1407788. Q. Ismail is funded by the College of Computer and Information Sciences, King Saud University, Riyadh, Saudi Arabia.

Authors’ addresses: R. Hoyle, Oberlin College, 10 N. Professor St, Oberlin, OH 44074; email: rhoyle@oberlin.edu; L. Stark, University of Western Ontario, London ON; email: cstark23@uwo.ca; Q. Ismail, King Saud University, Riyadh, Saudi Arabia; email: qalsmail@ksu.edu.sa; D. Crandall and A. Kapadia, Indiana University Bloomington, Bloomington, IN 47405; emails: djcran@indiana.edu, kapadia@indiana.edu; D. Anthony, University of Michigan, Ann Arbor, MI 48109; email: deniseum@umich.edu. 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.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1073-0516/2020/08-ART30 \$15.00

<https://doi.org/10.1145/3380960>

of content such as meal selections, “selfies,”¹ and even surreptitiously captured images from private life. Changes in information and communications technology (ICT) have enabled widespread sharing of personal photos. In 2015, Google users alone uploaded 13.7 petabytes worth of pictures to its Photos app.² A 2013 study found that 54% of Internet users posted original photos or videos online [18]. Photo-sharing is one of the most common activities of over two-thirds of American adults who now use social media [16, 72].³

The window on to our personal lives through posting or viewing photos of previously private activities, whether intimate or mundane, may affect not only individuals’ own preferences for sharing/viewing photos, but also how people view the appropriateness of such behavior (e.g., should it or should it not occur). Widespread sharing of personal photos online, in other words, is likely to shape social norms about privacy, i.e., commonly shared expectations about appropriate accessibility [6]. According to Helen Nissenbaum [60], human privacy entails a measure of “contextual integrity,” in that the flow of information, such as sharing personal photos, conforms to social norms about what is appropriate given particular social conditions. For example, it may be widely acceptable to share photos of a friend dressed up to go to a party, but sharing photos of a friend reading in her bedroom may be considered inappropriate either by the friend or by others.

Early studies of online photo-sharing explored users’ privacy preferences and found that the location of the photo, as well as user attitudes about privacy, influences whether images are shared or not [3, 85]. Similarly, user studies of wearable cameras (e.g., Google’s new Clips camera⁴) suggest that images with content exposing ostensibly private information (e.g., a visible computer monitor), or private spaces (e.g., home versus work), are kept more private (i.e., shared less) [32, 33]. Other aspects of social context, such as the number of people in the photo, also appear to influence whether lifeloggers share photos [33]. Much as how the invention of the automatic camera led to an explosion in photography in the late-19th-century America, which “radically altered the experience of seeing and being seen by others” [34, 43], we consider the social dimensions of privacy around personal photos online in the 21st-century United States. By examining shared privacy norms as well as individual privacy concerns and preferences, we can gain insight into the broader social dimensions of privacy and how new technologies affect fundamental patterns of social behavior. These second-order effects of technology and privacy [6, 21] can have unexpected and unintended consequences for the boundaries between what is considered private and public in a given society over time [34, 90].

As technologies change and new behaviors such as the widespread sharing of personal photos via social media become more common, existing privacy norms may change, and new norms may be created [6, 31, 60]. We conducted a large-scale online vignette study of perceptions of privacy of personal photos posted online to determine whether norms currently exist around the privacy of personal images, and if so, whether they vary based on social context characteristics. We examine basic features of social context in personal images of homes posted online, as well as the user’s relation to the image, to observe whether expectations of privacy are socially shared. While controlling for individual demographic characteristics, we also consider the role of personal privacy preferences and the various metrics for measuring them, seeking through this study to understand the role of both individual preferences and social norms in perceptions of privacy. By using

¹Oxford English Dictionary defines “selfies” as “a photograph one has taken of oneself, typically with a smartphone and shared via social media.” In 2014, 26% of Americans had shared a selfie: <http://www.pewresearch.org/fact-tank/2014/03/04/more-than-half-of-millennials-have-shared-a-selfie/>.

²<http://www.dailymail.co.uk/sciencetech/article-3619679/What-vain-bunch-really-24-billion-selfies-uploaded-Google-year.html>.

³<http://www.pewresearch.org/fact-tank/2017/01/12/evolution-of-technology/>.

⁴https://store.google.com/us/product/google_clips.

this type of experimental design, we can show whether and how particular features of homes or other aspects of social context differentially affect perceptions of privacy, both within and between subjects.

2 BACKGROUND—RELATED WORK

2.1 Privacy in Social Context: Privacy Norms

The concept of privacy varies across cultures and over time [5, 7, 17, 55, 68, 82, 91], and its particular definition further varies by academic discipline. We follow Anthony et al. [6] in defining privacy sociologically as the access of one social actor to another, in which access can vary along multiple dimensions, including “level” (high to low), “mode” (e.g., online versus face-to-face), and “type” (e.g., personal photos, health information). Much of the privacy scholarship examines individual privacy preferences: how much or little access (to self and others) individuals desire and value [1, 58, 83]. Other work explores how individuals vary in managing their own privacy, that is, their attempts to control access to, from, or about themselves or others (see, e.g., [5, 64, 83]). For example, the introduction of electronic medical records increased the privacy concerns of some patients, who then withheld information from their physicians as a way to (attempt to) manage privacy [11].

ICTs often blur previously established understandings and expectations of privacy because ICTs can radically change dimensions of accessibility [31, 46, 47, 58, 60, 71]. For example, changes in the mechanisms used by social media sites to disseminate information [9], or by utility companies for monitoring smart meters in homes [31], or by search engines to share personal information for commercial purposes [53], affect perceptions of privacy.

Accessibility is also shaped by social factors, including culture, laws, and social norms [1, 6, 54, 58]. Laws define the aspects of access that are legal and illegal [60, 68, 73, 91]. For example, in their landmark 1898 work on the legal right to privacy, Samuel Warren and Louis Brandeis cited the invention of portable cameras as a threat to individual privacy (particularly for the wealthy [80]). In 1902, another landmark legal case in the New York state involving the use of a portrait of a young woman for advertising purposes without her knowledge or consent led to statutory safeguards protecting the right to the private use of one’s own image in the commercial sphere [73].

In contrast to laws, social norms are informal, commonly held expectations about acceptable behavior that are socially enforced [28, 30].⁵ Shared expectations about appropriate levels or forms of accessibility to self and others, such as what are acceptable modes, levels, or types of access, are privacy norms [6, 30, 58, 60]. Nissenbaum [59, 60] argues that ICTs disrupt the contextual integrity of privacy because they affect privacy norms by changing the flows of information expected and considered appropriate in specific social contexts (see also [31, 53]). Since social norms are always being “created, altered, or negotiated” through behavior and interaction [20, 28], changes in technology that affect accessibility via information flows can cause the disruption of established privacy norms or the generation of new ones [31, 60].

In addition to being socially dynamic, norms are conditional, such that the same behavior can have different normative expectations depending on the social context [20, 28]. This variability means that privacy norms will not necessarily be uniform for a given technology or type of information, but instead may vary with other dimensions of accessibility. For example, Horne et al. [31] found that normative expectations related to smart-meter technology varied depending on the extent to which they allowed the utility companies to intervene directly in home appliances. Martin and Nissenbaum [53] show that expectations of privacy vary not only by type of information, but

⁵Social psychologists use the term injunctive norms to refer to the “ought-ness” of social norms, and distinguish this definition from the concept of descriptive norms, which refer to what most people do [13]. In this article, and consistent with sociological uses of the term [30], all references to norms are to injunctive norms.

also by recipient and use of information (see also [65]). Related work finds that privacy behavior is also related to contextual cues about the behavior of others [2, 39].

2.2 Privacy Preferences

Classic research on privacy preferences, such as from Westin [82, 83], measures privacy with survey questions to create typologies categorizing individuals based on how much accessibility to, from, or about themselves they think is important, and how strongly they value control over accessibility. Others scholars have developed more extensive typologies that build indices based on multiple dimensions of privacy attitudes. For example, Jensen et al. [38] propose an update to Westin's scale that focuses on online privacy specifically. Malhotra et al.'s Internet Users' Information Privacy Concerns (IUIPC) scale measures an "individual's subjective views of fairness within the context of information privacy" based on dimensions of awareness of information practices, concerns about collection of information, and desire for control over information [50]. The IUIPC has been used to show how individual privacy preferences are related to online behavior and use of technologies [50], while Jensen et al. and others question the value and validity of such scales and typologies (e.g., [38, 53]), particularly for explaining behavior [8].

Regardless of measurements, previous research shows that individuals do vary in their individual privacy preferences for sharing personal images online. Ahern et al. [3] explored image-sharing online through an interview study about Flickr images, finding that social identity concerns, as well as perceptions of security and convenience, affected user sharing behavior. Wu and Zhang [85] found that users varied in geotagging photos based on the type of location, geotagging less often for private locations. Other work on photo-sharing via lifelogging [32, 33] indicates that users actively attempt to manage their (and others') privacy when possible through a range of strategies for restricting image-sharing. In seeking to understand privacy norms for sharing personal images online, we therefore must also take into account the variation in individual privacy preferences.

2.3 Privacy and Sharing Digital Images

Previous studies have examined not only when and how digital photographs are shared online [3, 85], but also how to enable users to better manage their own privacy in the context of social norms around sharing digital images. Some research proposes mechanisms to obfuscate any shared image [77]. Choi et al. [12] suggest that editing photos (e.g., enhancing, cropping) can sometimes subvert algorithmic attempts to identify the image location. Similarly, DARKLY is a system layer developed by Jana et al. [35] that uses computer vision techniques to replace image contents with an "opaque reference" while letting apps access more private versions of the image. Other algorithmic tools evaluate a user's prerecorded preferences and take appropriate action when presented with images. Prasad et al. [65] analyzed how the privacy implications of wearable devices necessitated dynamic rather than preset preference settings. Toubiana et al. [25] argue that geolocation information should be leveraged to automatically apply preset privacy preferences when a photo is taken, while Klemperer et al. [41] contend that user-generated text tags can also be leveraged to guide automated privacy and access controls for online images. Fang and LeFevre [18] attempted to crowdsource privacy policies from a user's friends.

Other work has developed techniques to blur particularly sensitive parts of images. For example, Thomaz et al. [76] proposed rules that can be used to blur sensitive parts in images taken by lifeloggers in the context of logging eating behaviors. Korayem et al. [42] used computer vision techniques to detect monitors in lifelogging images and block them without the need of users to manually flag such images. Yu et al. [87–89] use deep machine learning to identify privacy-sensitive objects in shared images in order to apply a tool to automatically blur the privacy-sensitive content. Researchers have further studied the impact on the viewer's experience when transforming

parts of images for enhanced privacy and the trade off between privacy and utility for various filters across different scenarios [26, 27, 45]. Nonetheless, scholars as yet have a poor grasp of what people perceive as private for the purposes of image-sharing, possibly because privacy norms related to this behavior are in flux. A better understanding of privacy norms around personal images may enable the development of technical applications that recognize and protect image content while fulfilling users' expectations of privacy.

3 RESEARCH QUESTIONS

Prior research has focused on individual privacy preferences and the behavior for sharing images online. This work typically focuses on classifying individuals into categorical groups of users with high or low preferences for privacy (e.g., Westin's classic typology) and then observing how individuals (seek to) manage their own privacy. While we agree that individual privacy preferences and privacy management behavior are important, we build on the ideas of contextual integrity [53, 60] and sociology of privacy [6] to argue that privacy norms also play an important role in both perceptions and behavior related to privacy, but are much less understood. Here we seek to answer two research questions. (1) Do people share expectations about the privacy of contextual features of personal images of homes shared online? (2) (How) do perceptions of online-image privacy vary individually, based on the individual's (i) privacy preferences, and/or (ii) relationship to the image? We use contextual integrity and sociology of privacy to consider how features in images of the home may be perceived as more or less private.

As noted in Igo's history of privacy in the United States [34], the home has always been considered a significant and important site of privacy. Sociological scholars have described the extent to which privacy expectations vary within and across homes, such as in their size and layout [49], or between so-called "front-stage" public rooms (e.g., living and dining rooms) and "back-stage" private rooms (e.g., bedrooms) [23]. More recent research indicates that private areas of the home are considered "sensitive" information [44], and that particular social groups, such as college students, take active steps to manage access to such sensitive locales [67]. Prior research on wearable cameras [32, 33] has shown that images of bedrooms were less likely to be shared by lifeloggers, compared to other household rooms. Similarly, images containing computer monitors/screens (including smartphone screens) were defined as sensitive [44], and less likely to be shared than images without screens [32, 33], though this result does not hold in all studies [66]. Given that private information may be visible on computer monitors, this finding is consistent with ideas from the sociology of privacy and contextual integrity that images entailing greater access to personal information, such as what may be visible on a computer screen, will be considered more private than images with less access. Similarly, expectations of privacy are expected to be higher for images of intimate personal spaces, such as bedrooms, compared to other more public rooms in homes.

Based on these findings, we hypothesize the following:

- H1** Images of bedrooms will be perceived as more private/less appropriate to share online than images of other household rooms (kitchens, dining rooms, living rooms).
- H2** Images containing visible computer monitors/screens will be perceived as more private/less appropriate to share online than images without screens.

In addition to these features, we hypothesize that the presence of people in photos inside homes will be related to expectations of privacy. It is worth remembering that the first instant cameras and photos of people spurred some of the earliest writing expressing concerns about privacy, including Warren and Brandeis' famous Harvard Law Review article in 1890 [80]. Photos of a person provide access to a great deal of information about them, including not only about personal appearance (e.g., hair color, height), but also potentially about their location (via geotags and type of setting),

behavior, level of resources, and social network (via other people in the photo). Today, we know that people regularly remove tags of themselves in images posted on social media [48, 85]. In addition, studies of lifeloggers suggest that the number of people in an image may affect whether or not it is shared [32, 33, 66]. Given the extent of access provided by a photo of a person, we expect expectations about image privacy to vary based on whether it contains people and hypothesize that photos with people will be considered more private than images with no people.

H3 Images with people (one or more) will be perceived as more private/less appropriate to share online than images with no people.

By this logic, it may be the case that more people compared to fewer in a photo should be considered even more private. That is, perceptions of image privacy will simply increase with the number of people in the image. However, more people may indicate a more social, even public, versus private setting or situation, so it is not clear that images with more people will be perceived as more private than images with fewer or no people. Below we explore whether images with two people compared to none or one are considered more or less private as an open research question rather than a specified hypothesis. We discuss the implications and additional research questions in Sections 5.4 and 6.

As noted above, though we seek to identify privacy norms for image-sharing, we also expect individual preferences for privacy to affect perceptions of privacy for personal images shared online. Many studies have found associations between individuals' privacy attitudes and preferences and their information-sharing intentions and behavior (see Belanger and Crossler [8]). A number of different attitudinal measures of privacy preferences exist and no one measure is considered to be definitive. Thus, we examine multiple measures of privacy preferences. Though there has long been a concern and criticism of such attitudinal scales [38, 50], particularly for explaining individual privacy behavior [24, 84], such attitudes are expected to affect perceptions of privacy, such that the following hypothesis holds:

H4 Individuals with stronger preferences for privacy will be more likely to perceive images as private/less appropriate to share online.

In addition to individual privacy preferences as measured through privacy attitudes and concerns, we explore how the participant's relation to the image may be related to perceptions of privacy and thus shed further light on privacy norms. In the current context of social media, people have the opportunity to view the images of others, to post images they have taken, and to have images of them posted online (whether by themselves or by others). Thus, online photo-sharing has implications for perceptions of privacy related to oneself and to others, some of whom may be merely "bystanders," that is, those who are (knowingly or unknowingly) captured in an image (or by any type of sensing device) [63]. We know that bystanders express concerns about being captured by sensing devices [15, 86], and are emotionally and viscerally sensitive to their own privacy [74]. Studies also find that people express concerns about the privacy of others, and take steps to protect the privacy of bystanders [4, 32, 33]. Still, we expect that images that entail more access to the participant (e.g., participant visualizes being in image) will be perceived as more private than images viewed or posted online.

H5 Expectations of image privacy will vary by participant's relation to the image such that images of the participant or of the participant's home will be perceived as more private/less appropriate to share than (i) images taken by participants, and (ii) images viewed online by participants.

Table 1. Study Conditions by Image Characteristics

Condition	Type of Household Room	Monitor/Screen	Number of People
1	Bedroom	None	0
2	Bedroom	Yes	0
3	Other room	None	1
4	Bedroom	None	1
5	Other room	Yes	1
6	Bedroom	Yes	1
7	Other room	None	2
8	Bedroom	None	2
9	Other room	Yes	2
10	Bedroom	Yes	2

Note: Other room = kitchen, living room, or office.

4 STUDY PROCEDURES

4.1 Study Design

We conducted an online factorial vignette study with 279 Amazon’s Mechanical Turk (mTurk) workers in order to examine privacy perceptions about personal images shared online. Vignette studies use an experimental design in which hypothetical stories/scenarios, with details corresponding to experimentally defined conditions, are deployed to test for the variation in judgments either between or within subjects [36, 37]. Vignette studies have been used to explore privacy issues [31, 53, 56] by employing scenarios to explore subjects’ perceptions and expectations about sharing or exposing different types of information under different conditions. Here, the experimental conditions are features in images that vary by the hypothesized characteristics described above (i.e., bedrooms versus other household rooms; monitors present or not; zero, one, or two people). For each image-vignette, subjects were asked questions about the content of the image (to confirm their view of the image conformed exactly to the experimental condition), as well as about their judgments of whether it was private and thus not appropriate to be posted online for each of the following three scenarios: (i) if they saw the photo posted online; (ii) if they took the photo; and (iii) if they were in the photo (or for images with no people, if the room showed was in their home). In addition, after viewing all the image-vignettes, subjects were asked about their online behavior and general privacy attitudes and behaviors; prior to the vignettes, subjects answered basic demographic questions (see Online Appendix A for the survey questions and an example of image-vignettes). Each subject was presented with 10 “image-vignettes” consisting of an image that varied across conditions of type of household room (bedrooms versus other rooms, including kitchens, living rooms, or offices), whether a computer monitor/screen (including smartphone screens) was present or not, and the number of people (zero, one, or 2) (see Table 1). The order of vignettes was randomized for each subject. The study design and vignette-survey materials were approved by our university’s institutional review board.

4.2 Selection of Images and Respondent Sample

For each condition, we selected six images from the publicly available online image-sharing sites Flickr and Google Images, searching for images with terms such as “one person in bedroom.”⁶

⁶No photos are included with people in any stage of undress, or in any activity that might be considered “sexual” or intimate.

The figure displays two side-by-side screenshots from a survey instrument. The left screenshot shows a photograph of a room with a desk, chair, and bed. Below the photo is a question: "What type of room do you think this photo depicts?" followed by a list of radio button options: Bedroom, Living Room, Kitchen, Bathroom, Office, and Other (Please specify):. The right screenshot shows a text prompt: "If you saw this photo on social media, even if you didn't know the place, how strongly do you agree or disagree with the following statement?" followed by the statement: "I think it is private and should **not** be posted to social media." Below this is a set of five radio button options: Strongly agree, Agree, Neither Agree nor Disagree, Disagree, and Strongly Disagree.

Fig. 1. Sample screenshots of the survey instrument.

We conducted a pilot study with 10 subjects to ensure all 60 photos (10 conditions $\times$ 6 images per condition) captured the condition-specific characteristics and subjects correctly identified the experimental features of each image. Images that did not match 100% for any of the pilot subjects were removed and new photos added as needed until our pilot testers showed agreement on the condition-specific features for all six images per condition. In addition, we determined that the vignette questions (described in Section 4.3) were understood by all pilot respondents.

A scope condition for this study was to constrain the identifiable characteristics of the people shown in the images containing them, in an attempt to limit confounding elements as much as possible. We therefore sought to ensure images conformed to the hypothesized conditions for number and apparent gender of individuals, while controlling age (adults only) and race/ethnicity (apparently White/Caucasian only). Setting this condition was not because we expect the factors of age and race/ethnicity do not matter; in contrast, there is clear evidence that privacy is unequally distributed, including by age and race/ethnicity [6, 10]. We note the implications for these limitations in Sections 6 and 7.

During the study, subjects were asked to first identify the features in each image, i.e., the type of room, whether a monitor was visible, and the number of people in the image (see Figure 1 for an example vignette-image and questions; see also Online Appendix A, where study subjects were randomized between the six images within each condition, and also randomized as to the order in which they were presented each condition).

Respondents were selected through Amazon's mTurk platform, targeting "Master" surveyors⁷ with the description "Answer questions about what is contained in various images, and your feelings about them. The study is being conducted through Indiana University. You will be compensated \$2.25 for taking the survey." Respondents were given 1 hour to complete the survey, though our initial pilot estimated that it should take 20 minutes; thus, payment was at the rate of \$7 per hour for compensation. Six attention-check questions that simply asked respondents to choose one of the response categories listed were included within the survey. Respondents were paid if they successfully answered more than three of the attention questions, but their data were discarded if they failed to answer any of the attention questions. A total of 418 respondents completed the survey, with 279 (66%) completing all six of the attention-check questions correctly. Recently (but

⁷"Master" surveyors are those workers that have a proven track record at successfully completing a variety of tasks.

after we conducted the study), research has suggested that attention questions should not be used in surveys or online studies [79]. This research finds that attention questions, like the ones we used to ensure that study respondents were attending to the questions and not simply clicking on responses randomly, can cause a selection effect in which those who fail attention questions are systematically different from those who do not. We went back to the original data and compared demographic characteristics between those we excluded for failing the attention questions to the respondents in the analytic sample. We found no differences by gender, age, education, or race between those excluded and the final analytic sample, suggesting that the attention-check questions did not have a biasing effect on our sample.

4.3 Measures

For each image-vignette, subjects were asked to consider the following questions, all with response categories on a 5-point Likert scale of 5 = strongly agree private to 1 = strongly disagree private:

- (1) If you saw this photo on social media, even if you didn't know the people, how strongly would you agree or disagree with the following statement? I think it is private and should not be posted to social media.
- (2) If you were the person who took this photo, how strongly would you agree or disagree with the following statement? I think it is private and should not be posted to social media.
- (3a) (For photos without people): If this photo were of your home, how strongly would you agree or disagree with the following statement? I think it is private and should not be posted to social media.
- (3b) (For photos with people): If you were the/a person in this photo, how strongly would you agree or disagree with the following statement? I think it is private and should not be posted to social media.

We created three dichotomous dependent variables (saw photo, took photo, and own-home/in-photo) in which responses of "strongly agree private" and "agree private" are coded as 1, and all other responses coded as 0 for each of the questions above, collapsing own-home/in-photo responses into a single variable.

The key independent variables for measuring social context factors within images are the categories: type of household room (bedroom = 1, else = 0), presence of computer monitors/screens (including smartphone screens = 1, no screen = 0), and number of people in the photo (none, **1Male**, **1Female**, **2Females**, **1Male&1Female**).⁸

In addition to the influence of social context factors in the image, we also hypothesized that individual privacy preferences will influence expectations of privacy for personal images shared online. Though no agreed-upon measures of individual privacy preferences exist, we used four different measures (defined below). The components of each privacy preference measure were asked after the experimental conditions as statements in randomized order.

The best known measure of privacy preferences is probably Westin's classic privacy scale [82, 83], which we include using Westin's classic three questions, all with response categories on a 5-point Likert scale of 5 = strongly agree to 1 = strongly disagree:

- Consumers have lost all control over how personal information is collected and used by companies.

⁸Images with two males were cut from the final analysis because of limited data.

- Most businesses handle the personal information they collect about consumers in a proper and confidential way. (Recoded response so high value = more privacy concern)
- Existing laws and organizational practices provide a reasonable level of protection for consumer privacy today. (Recoded response so high value = more privacy concern)

Subjects giving privacy-oriented answers (responses = 4 or 5) to all three questions are classified, according to the original typology, as Westin-fundamentalists, those giving no privacy-oriented answers are classified as Westin-unconcerned, while those in-between are classified as Westin-pragmatists.

Jensen et al. [38] proposed an update of Westin's scale based on five privacy statements more directly related to online activities. Though Jensen et al. [38] recommend creating the same three categories as Westin's original framework, here we create a mean of the five statements, all with 5-point Likert response categories:

- I am concerned about online identity theft.
- I am concerned about my privacy online.
- I am concerned about my privacy in everyday life.
- I am likely to read the privacy policy of an ecommerce site before buying anything. (Reverse coded)
- Privacy policies accurately reflect what companies do. (Reverse coded)

Another attitudinal measure of privacy preferences, proposed by Malhotra et al. [50] is the IUIPC scale, which is part of their larger behavioral model for online behavior. The IUIPC includes three subscales on privacy "awareness," "collection," and "control," all with 5-point Likert response categories⁹:

Awareness:

- Companies seeking information online should disclose the way the data are collected, processed, and used.
- A good consumer online privacy policy should have a clear and conspicuous disclosure.
- It is very important to me that I am aware and knowledgeable about how my personal information will be used.

Collection:

- It usually bothers me when online companies ask me for personal information.
- When online companies ask me for personal information, I sometimes think twice before providing it.
- It bothers me to give personal information to so many online companies.
- I'm concerned that online companies are collecting too much personal information about me.

Control:

- Consumer online privacy is really a matter of consumers' right to exercise control and autonomy over decisions about how their information is collected, used, and shared.
- Consumer control of personal information lies at the heart of consumer privacy.
- I believe that online privacy is invaded when control is lost or unwillingly reduced as a result of a marketing transaction.

⁹These authors use a 7-point Likert response scale in the original model [50].

Table 2. Descriptive Statistics of Subject Characteristics, $n = 279$

Variables	%
Female	56%
Non-White Race	19%
Age	
18–29 years of age	24%
30–39 years	37%
40–49 years	18%
50–59 years	13%
60+ years	8%
Education Level	
% with College degree or higher	65%
Use Social Media	
several times/day (=1) vs less often (=0)	53%
Use Mobile Phone	
several times/day (=1) vs less often (=0)	81%
Use Phone for Photo-Sharing	
few times/month or more (=1) vs less often (=0)	57%

We computed the mean score for each subscale to create three variables of **IUIPC-awareness**, **IUIPC-collection**, and **IUIPC-control**.¹⁰

Finally, we create our own simple self-rated privacy question: “Which of these statements more accurately describes you: I am generally a private person and like to keep to myself. OR I am generally an open person who enjoys sharing with others.”

In all analyses, we control for subject demographics (gender, race, education, and age) and the frequency of social media and mobile phone use, coded as 1 = daily use and 0 = less than daily use, as well as the frequency of photo-sharing via the phone, coded as 1 = few times/month or more, and 0 = less often. Table 2 shows the descriptive statistics for all subject characteristics (discussed further in Section 5.1).

4.4 Method of Analysis

This study employs a within-subject study design to examine privacy norms related to features of images shared online, while also analyzing whether between-subject differences in privacy preferences also influence privacy perceptions. We use an extended generalized linear mixed-model (GLMM) that allows for the decomposition of within- and between-subject effects for nonlinear outcomes using the `xthybrid` command with a logit-link function in Stata-13 [69, 70]. Though fixed-effects models are useful for repeated-measures designs like ours, enabling estimates of within-subject variables while adjusting for subject-invariant variation, they cannot provide

¹⁰Malhotra et al. [50] propose using principal component factor analysis to create one latent factor with the three subscales as component measures of it. In separate analyses (not shown), we calculate a one-factor IUIPC measure ($\alpha = 0.86$), and find that using the one-factor IUIPC measure instead of the three reported here produces results substantively identical to those reported later, though it is never statistically significant.

effect estimates for the between-subject variables. Hybrid statistical models provide estimates of both within-subject effects, that is, how on average a within-subject change in an independent variable (e.g., room type) is associated with a within-subject change in the dependent variable (expectation of privacy), as well as between-subject effects, that is, how individual differences in privacy preferences are associated with differences in the dependent variable (i.e., expectation of privacy). This is accomplished by (i) specifying subject-varying independent variables as deviations from the subject (cluster) mean, and (ii) adding the means of the original subject-varying variables to the model. Thus, we can evaluate whether there are between-subject effects (by privacy preferences) in addition to the within-subject effects (of image features) on perceptions of privacy.

For each dependent variable, we examine a series of separate models that include the within-subject image features and the different measures of privacy preferences. All models calculate robust standard errors adjusted for the clustering of observations within subjects. We report odds ratios (OR) with 95% confidence intervals for each of the hypothesized within-subject and between-subject variables, while controlling for individual-level (between-subject) fixed effects of demographic characteristics and technology use. In these models, the OR represent the odds that the proportion of respondents rating the photo as private is higher (or lower) given the presence of the particular variable (e.g., bedroom, one person) compared to the odds of the proportion rating the photo as private in the absence of the variable (i.e., not a bedroom, no people). We used a significance threshold of 0.05 to determine whether or not a variable was significant. Based on our sample of 279 subjects with repeated measures across up to 10 observations each (mean number of observations per subject = 6.5), and mean proportions rating the image as private across the dependent variables of SAW photo = 0.185, TOOK photo = 0.192, and IN photo = 0.259, we estimate 80%–99% power with effect sizes of .05 to .09. These effect sizes indicate that for a hypothesis that a photo feature will increase the likelihood of an image being rated as private, where the mean proportion is .19, we have 80% power if the proportion rating as private is .24, and 99% power if the proportion is .28.

In addition to these models, we conducted additional analyses (not shown here) to evaluate the robustness of the findings and discuss them briefly below.

5 STUDY RESULTS AND DISCUSSION

5.1 Subject Characteristics

A total of 279 subjects completed all aspects of the study via mTurk, including correctly responding to all six attention-check questions and correctly identifying the key features in each photo condition. Table 2 shows that 56% of the sample is female, and 19% were of non-White race (self-rated as one of following: African American/Black, Asian American/Pacific Islander, Hispanic; we collapse these categories into one for analysis). The majority of subjects are younger than 39 years of age, with about 20% over the age of 50 years. Consistent with the profile of master Turk users, this is a highly educated and technology-savvy sample, with two-thirds of the subjects having at least a college education. Most use a mobile phone and social media several times a day, as well as sending/receiving photos on their phones at least a few times per month. Given our study is designed to test both within- and between-subject conditions rather than provide population estimates of attitudes or behavior, we did not use a random sample of the population. We discuss how the sample characteristics, and the use of mTurk as a tool for research more generally, may be relevant for our findings in Section 7.

Table 3 shows descriptive statistics for the measures of individual privacy preferences. According to Westin's scale, 24% of the sample are privacy fundamentalists, 26% are privacy pragmatists,

Table 3. Descriptive Statistics of Individual Privacy Preference Measures, $n = 279$

Privacy Preference Measure	% or Mean (standard deviation)
Westin Privacy Categories	
Privacy Fundamentalist	24%
Privacy Pragmatist	26%
Privacy Unconcerned	50%
Jensen Online Privacy Mean	3.6 (.54) Range = 1.8, 4.8
Internet Users' Information Privacy Concerns (IUIPC) Scale	
IUIPC-Awareness mean	4.4 (0.01)
IUIPC-Collection mean	4.0 (0.01)
IUIPC-Control mean	4.1 (0.01)
Self-Rated Private (private person=1, open person=0)	83%

and 50% are privacy unconcerned. For the 5 privacy statements recommended by Jensen et al. [38], the mean is 3.6 on a 5-point Likert scale (higher scores = more private), which indicates moderate to high privacy concerns on average, with a range (1.8–4.8) including both low and high privacy scores. The Malhotra et al. [50] IUIPC awareness, collection, and control mean scores of at or above 4 on a 5-point Likert scale (higher scores = more private) indicate relatively high levels of privacy concerns in each of the three areas. Finally, the vast majority of the subjects (83%) self-rate themselves as a “private person” rather than an “open person.”

5.2 Findings: Privacy Norms for Image Features

Tables 4–6 show the results of hybrid mixed-effects GLM models with logit-link functions for each of the three dependent variables: expectations of privacy if the subject saw the photo posted online (Table 4), took the photo (Table 5), or if the photo was of their home/they were in the photo (Table 6). Each table includes five separate models to test each of the different privacy preference measures. Model 1 in each table has no individual privacy preferences, model 2 includes the Westin privacy measures, model 3 includes the Jensen et al. [38] online privacy mean, model 4 includes the IUIPC privacy measures [50], and model 5 includes our own measure of self-rated privacy. All models also include controls for subject gender, race, age, education, and frequencies of mobile phone use, social media use, and phone photo-sharing.¹¹ Each model includes only observations (number of subjects multiplied by the number of responses to each image question) with nonmissing data on all variables, so the number of subjects (the cluster variable in hybrid mixed-effects models) and the number of observations (within-subject observations) vary slightly between models.

To evaluate the first hypothesis—that images of bedrooms will be considered more private than images of other household rooms—we can look at the results in row 1 of each table. Across each of the dependent variables shown in Tables 4–6, the OR for bedroom are never statistically significant, and we thus reject hypothesis H1.

¹¹Note that the hybrid mixed models allow us to include demographics as between-subject fixed effects. However, none of the demographic characteristics are significant in any of the models.

Table 4. OR [with 95% Confidence Intervals] for Expectations of Privacy if SAW Photo Posted Online

Image Characteristics	Model 1: Baseline	Model 2: Westin Privacy	Model 3: Jensen Online Privacy	Model 4: IUIPC Means	Model 5: Self-Rated Privacy
WITHIN Subjects:					
Bedroom ¹	0.84 [.58–1.2]	0.84 [.58–1.2]	0.85 [.59–1.2]	0.87 [.60–1.3]	0.84 [.58–1.2]
Monitor / Screen	0.68** [.51–.91]	0.68** [.51–.91]	0.68** [.51–.91]	0.69** [.51–.92]	0.68** [.51–.91]
Number of People²					
One Male	2.3*** [1.4–3.7]	2.3*** [1.4–3.7]	2.3*** [1.5–3.7]	2.3*** [1.5–3.7]	2.3*** [1.4–3.7]
One Female	3.7*** [2.3–5.5]	3.7*** [2.3–5.5]	3.7*** [2.3–5.5]	3.7*** [2.3–5.5]	3.7*** [2.3–5.5]
Two Females	0.71 [.44–1.1]	0.71 [.44–1.1]	0.71 [.44–1.1]	0.71 [.44–1.1]	0.71 [.44–1.1]
Two: Male & Female	0.15** [.05–.54]	0.15** [.05–.54]	0.15** [.05–.54]	0.15** [.05–.54]	0.15** [.05–.54]
BETWEEN Subjects: Privacy Preferences					
Westin Privacy³					
Fundamentalists	—	1.3 [.82–2.2]	—	—	—
Unconcerned	—	1.2 [.73–1.9]	—	—	—
Jensen Online Privacy					
IUIPC Scale	—	—	1.5 [.95–2.4]	—	—
Awareness	—	—	—	0.89 [.48–1.6]	—
Collection	—	—	—	1.2 [.81–1.6]	—
Control	—	—	—	1.1 [.73–1.8]	—
Self-Rated Privacy					
Constant	0.12 [.01–2.1]	0.12 [.01–2.1]	0.02* [.001–.68]	0.07 [.001–4.1]	2.8*** [1.5–4.9]
	–LL = –761.2	–LL = –760.5	–LL = –756.9	–LL = –754.1	–LL = –755.4
	Wald $\chi^2 = 147.6^{***}$	Wald $\chi^2 = 155.9^{***}$	Wald $\chi^2 = 145.6^{***}$	Wald $\chi^2 = 146.8^{***}$	Wald $\chi^2 = 157.3^{***}$
	df = 22	df = 24	df = 23	df = 25	df = 23
	No. clusters = 274	No. clusters = 274	No. clusters = 273	No. clusters = 270	No. clusters = 274
	No. obs = 1,780	No. obs = 1,780	No. obs = 1,774	No. obs = 1,755	No. obs = 1,780

Notes: * $p < .05$; ** $p < .01$; *** $p < .001$. ¹Versus other rooms; ²versus zero people; ³versus Westin privacy-pragmatists.

Each model uses a hybrid mixed-effects GLM model (with logit-link function), and includes covariate controls for gender, race, age, education, frequencies of mobile phone use, social media use, and phone photo-sharing.

Table 5. OR [with 95% Confidence Intervals] for Expectations of Privacy if TOOK Photo Posted Online

Image characteristics	Model 1: Baseline	Model 2: Westin Privacy	Model 3: Jensen Online Privacy	Model 4: IUIPC Means	Model 5: Self-Rated Privacy
WITHIN Subjects:					
Bedroom ¹	0.96 [.65–1.4]	0.96 [.65–1.4]	0.98 [.66–1.5]	0.99 [.68–1.5]	0.96 [.65–1.4]
Monitor/Screen	0.68** [.52–.89]	0.68** [.52–.89]	0.68** [.52–.89]	0.68** [.52–.89]	0.68** [.52–.89]
Number of People²					
One Male	2.8*** [1.8–4.5]	2.8*** [1.8–4.5]	2.9*** [1.8–4.5]	2.9*** [1.8–4.9]	2.8*** [1.8–4.5]
One Female	3.9*** [2.6–6.1]	3.9*** [2.6–6.1]	3.9*** [2.6–6.1]	3.9*** [2.6–6.1]	3.9*** [2.6–6.1]
Two Females	0.83 [.53–1.3]	0.83 [.53–1.3]	0.83 [.53–1.3]	0.83 [.53–1.3]	0.83 [.53–1.3]
Two: Male & Female	0.22* [.06–.81]	0.22* [.06–.81]	0.22* [.06–.81]	0.22* [.06–.84]	0.22* [.06–.81]
BETWEEN Subjects: Privacy Preferences					
Westin Privacy³					
Fundamentalists	—	1.3 [.80–2.0]	—	—	—
Unconcerned	—	1.1 [.69–1.9]	—	—	—
Jensen Online Privacy					
IUIPC Scale	—	—	1.3 [.84–2.0]	—	—
Awareness	—	—	—	.92 [.52–1.6]	—
Collection	—	—	—	1.3 [.92–1.8]	—
Control	—	—	—	0.85 [.57–1.3]	—
Self-Rated Privacy					
Constant	0.03* [.002–.68]	0.03* [.002–.66]	0.01** [.0004–.38]	0.04* [.001–1.6]	3.8*** [2.1–6.7]
	–LL = –784.8 Wald $\chi^2 = 145.3^{***}$ df = 22 No. clusters = 274 No. obs = 1,783	–LL = –784.3 Wald $\chi^2 = 149.8^{***}$ df = 24 No. clusters = 274 No. obs = 1,783	–LL = –781.6 Wald $\chi^2 = 142.7^{***}$ df = 23 No. clusters = 273 No. obs = 1,777	–LL = –776.7 Wald $\chi^2 = 144.6^{***}$ df = 25 No. clusters = 270 No. obs = 1,758	–LL = –774.2 Wald $\chi^2 = 158.6^{***}$ df = 23 No. clusters = 274 No. obs = 1,783

Notes: * $p < .05$; ** $p < .01$; *** $p < .001$. ¹Versus other rooms; ²versus zero people; ³versus Westin privacy-pragmatists.

Each model uses a hybrid mixed-effects GLM model (with logit-link function), and includes covariate controls for gender, race, age, education, frequencies of mobile phone use, social media use, and phone photo-sharing.

Table 6. OR [with 95% Confidence Intervals] for Expectations of Privacy if IN Photo or of OWN Home Posted Online

Image characteristics	Model 1: Baseline	Model 2: Westin Privacy	Model 3: Jensen Online Privacy	Model 4: IUIPC Means	Model 5: Self-rated Privacy
WITHIN Subjects:					
Bedroom ¹	0.89 [63–1.3]	0.89 [63–1.3]	0.89 [64–1.3]	0.91 [65–1.3]	0.89 [63–1.3]
Monitor/Screen	0.73* [56–96]	0.73* [56–96]	0.73* [56–96]	0.74* [57–97]	0.73* [56–96]
Number of People²					
One Male	2.1*** [1.4–3.2]	2.1*** [1.4–3.2]	2.1*** [1.4–3.2]	2.2*** [1.4–3.3]	2.1*** [1.4–3.2]
One Female	2.6*** [1.8–3.7]	2.6*** [1.8–3.7]	2.6*** [1.8–3.7]	2.6*** [1.8–3.7]	2.6*** [1.8–3.7]
Two Females	0.61** [41–89]	0.61** [41–89]	0.61** [41–89]	0.61** [41–89]	0.61** [41–89]
Two: Male & Female	0.07*** [02–30]	0.07*** [01–30]	0.07*** [02–30]	0.07*** [02–30]	0.07*** [02–30]
BETWEEN Subjects: Privacy Preferences					
Westin Privacy³					
Fundamentalists	–	1.3 [.84–1.9]	–	–	–
Unconcerned	–	1.1 [.70–1.8]	–	–	–
Jensen Online Privacy					
IUIPC Scale	–	–	1.7** [1.2–2.5]	–	–
Awareness	–	–	–	1.1 [.63–1.9]	–
Collection	–	–	–	1.4* [1.0–1.9]	–
Control	–	–	–	0.98 [.69–1.4]	–
Self-Rated Privacy					
Constant	0.19 [01–3.3]	0.19 [01–3.3]	0.02* [001–.58]	0.22 [001–1.8]	3.0*** [1.8–4.9]
	–LL = –914.3	–LL = –913.8	–LL = –906.9	–LL = –902.7	–LL = –905.5
	Wald $\chi^2 = 132.7^{***}$	Wald $\chi^2 = 133.2^{***}$	Wald $\chi^2 = 134.6^{***}$	Wald $\chi^2 = 138.2^{***}$	Wald $\chi^2 = 142.3^{***}$
	df = 22	df = 24	df = 23	df = 25	df = 23
	No. clusters = 274	No. clusters = 274	No. clusters = 273	No. clusters = 270	No. clusters = 274
	No. obs = 1,788	No. obs = 1,788	No. obs = 1,782	No. obs = 1,763	No. obs = 1,788

Notes: * $p < .05$; ** $p < .01$; *** $p < .001$. ¹Versus other rooms; ²versus zero people; ³versus Westin privacy-pragmatists.

Each model uses a hybrid mixed-effects GLM model (with logit-link function), and includes covariate controls for gender, race, age, education, frequencies of mobile phone use, social media use, and phone photo-sharing.

We can evaluate hypothesis H2 that images with computer monitors/screens will be considered more private than images with no screens, by looking at the OR reported in row 2 of each of the three tables: All are significant in all models for all three dependent variables—however, all are less than 1, indicating that images with computer screens are statistically less likely to be considered private than images with no screens, all else equal. This finding that images with computer screens have, on average, about 30% lower odds of being considered private is the opposite of hypothesis H2 and a somewhat surprising finding, given that computer screens often contain personal information that people may consider private. Not only does this result contradict hypothesis H2, it also suggests that expectations about the privacy of images containing computer screens may be more complicated than simply considering the possibility of personal information exposure. Since none of the computer screens in the images used in the study included any visible content, no information exposure was possible and this may be the case why screens were not considered private. Furthermore, the ubiquity of personal computers (whether desktops, laptops, or smartphones) in our personal lives and therefore also in our personal photos may be why they are actually considered less private, especially when no content is visible.

Hypothesis H3 considers the privacy norms related to the presence and number of (apparently White/Caucasian) people in a photo. As shown in row 3 in each of Tables 4–6, we can see that images with one male or with one female are always significant and positive, with OR indicating that they are two to almost four times as likely to be considered private compared to the reference category of zero people in the photo, supporting hypothesis H3. In contrast, images with two people, specifically those with a male and female, are consistently significantly less likely to be considered private, by about 80%–90%, than photos with no people across all three dependent variables. Similarly, images with two females are significantly less likely to be evaluated as private than images with no people, by about 40%, but only for scenarios in which the subject considers being in the photo (Table 6). Overall, these findings show that compared to an image with no people, photos with one single person are two to three times more likely to be evaluated as private, while photos with two people, particularly a male and female, are about half as likely to be evaluated as private. Support for hypothesis H3 indicates that privacy norms for images are strongly related to the social conditions, which in this case is the number of people, as depicted in the image. We explore these findings more in Section 5.4, and further discuss their implications in Sections 6 and 7.

Hypothesis H4 considers how individuals' personal privacy preferences also affect expectations of privacy, separately from privacy norms. Westin's classic three categories of privacy fundamentalists, privacy pragmatists (the suppressed category), and privacy unconcerned are tested in model 2 for each dependent variable (shown in column 2 in Tables 4–6); but, none is ever statistically significant. These findings provide evidence consistent with previous scholarship, suggesting that the Westin categories do not appear to be particularly useful in understanding privacy in online contexts.

The online privacy scale recommended by Jensen et al. [38] is tested in model 3 for each dependent variable. As with the Westin categories, the Jensen Online Privacy mean is not significant for privacy if the subject saw photo online (Table 4) or took the photo (Table 5), but it is significant and positive for subjects imagining being in the photo or as a photo of their own home (Table 6). This latter finding shows that individuals with greater privacy concerns, as measured by the Jensen Online Privacy mean, are about one-and-one-half times more likely to evaluate an image they are in as private, over and above the normative expectations of privacy related to the features of the image. This finding provides support for hypothesis H4 and suggests that the Jensen et al. scale for online privacy is a useful tool.

Model 4 for each dependent variable examines the effects of the Malhotra et al. [50] IUIPC scales for awareness, collection, and control. None of the IUIPC scales are significant for privacy

† Within all number-of-people categories except **1Male&1Female**, Repeated-measures ANOVA (adjusted for between-subject characteristics) with post-hoc Bonferroni tests show own home/in photo is significantly higher privacy than saw photo ($p < 0.001$) and took photo ($p < 0.001$). There is no statistical difference between saw photo and took photo in any category except **1Female**. Separately, Bonferroni tests show there are significant differences within the photo categories:

- Saw Photo: **1Male** and **1Female** images significantly more likely private ($p < 0.001$) than all other categories. No statistical difference between no-people and **2Females** categories. Category with **1Male&1Female** images significantly less likely private ($p < 0.001$) than all other categories.
- Took Photo: **1Male** and **1Female** images significantly more likely private ($p < 0.001$) than all other categories. No statistical difference between no-people and **2Females** categories. Category with **1Male&1Female** images significantly less likely private ($p < 0.05$) than all other categories.
- Own Home/In Photo: All categories are statistically different from all others ($p < 0.001$): **1Female** > **1Male** > no-people > **2Females** > **1Male&1Female**.

Fig. 2. Image privacy by the subject's relation to the photo and number of people in the image.

if the subject saw the photo online (Table 4) or took the photo (Table 5). For subjects imagining being in the photo or as a photo of their own home (Table 6), UIPC-collection is significant and positive, indicating those subjects who are more concerned about online information collection are more likely to perceive a photo as private and that it should not be posted online. This finding for the UIPC-collection scale is consistent with the finding for the Jensen Online Privacy mean, and provides some additional support for hypothesis H4.

Finally, model 5 shows a strongly significant positive effect of our dichotomous self-rated privacy measure for all three dependent variables, indicating that those who identify as “a private person” are two to nearly four times more likely to say that images are private and should not be posted online, even after controlling for the features of the image. Based on the Wald χ^2 for each model (a measure of model fit), model 5 shows the best fit for each dependent variable. This finding for self-rated privacy also supports hypothesis H4 that individual privacy preferences influence expectations of image privacy online.

To summarize findings related to hypotheses H1–H4, we find no support for hypothesis H1 that privacy expectations for images of bedrooms are greater than for other, more public household rooms. For hypothesis H2 regarding privacy norms for images showing a computer monitor/screen, the analyses show an unexpected finding that directly contradicts hypothesis H2: Such photos are significantly less likely to be evaluated as private than photos without monitors/screens. We find strong support for hypothesis H3 that images with one person are more likely to be considered private than images with no people. We also considered what the privacy expectations might be for two people in a photo and found that images with two people, particularly a male and a female, were significantly less likely to be considered private than photos with no people. Figure 2 illustrates the pattern of privacy expectations across the categories for the number of people in the photo, and shows that within each dependent variable (saw photo, took photo, own home/in photo), privacy expectations are highest for photos with one person (particularly one female) and lowest for two people (particularly a male and a female).

5.3 Findings: Subject Relation to Photo

Figure 1 also illustrates how subjects' relation to the photo, as measured through the three dependent variables, influences privacy expectations as defined in hypothesis H5, in which greater access increases expectations of privacy. Within each category for the number of people in the image, within-subject means for the percentage of images evaluated as private are significantly higher when the subject imagines being in the image or an image of their home compared to either seeing the photo posted online or taking the photo. Using repeated-measures ANOVA (adjusted for between-subject characteristics) with Bonferroni tests shows that among images of no people, subjects are more likely to evaluate the image as private if they considered it to be an image of their own home compared to seeing it posted online or taking it themselves [20.8% compared to 12.9% and 13.1%, respectively]. This pattern holds across all number-of-people image categories except those with two people consisting of a male and a female. In all number-of-people categories, there are no statistical differences in evaluations of privacy between saw photo or took photo. Overall, these findings support hypothesis H5 that subjects' role in relation to images posted online influences expectations of privacy; specifically, images that enable greater access to the subject are more likely perceived as private than those with lower/no access. These findings are consistent with theoretical ideas that people are viscerally sensitive to their own privacy [74].

5.4 Robustness Checks

We conducted a number of additional statistical tests (not shown) to evaluate the robustness of our findings. First, to evaluate whether the findings were influenced by specific images, we reran all models reported in Tables 4–6 after dropping outlier images (those with the highest and lowest mean privacy ratings in each condition). Our findings are substantively identical, including findings of statistically significant and nonsignificant relationships, to those reported in the tables.

Second, since our experimental design is not fully crossed, such that there is no condition with images of a nonbedroom with zero people, the independent variables for room type and number of people are correlated. Therefore, we examine key variables (type of room, monitor presence) within each of the conditions for the number of people above zero (1–2), as well as the reverse (number of people within each room/monitor condition), using the covariates as in the self-rated privacy model (#5) for each dependent variable shown in Tables 4–6. We find that all substantive and statistically significant relationships reported in Tables 4–6 hold in these subgroup analyses, indicating that our reported findings for the independent effects of room type, monitor presence, and number of people are robust.

Third, in considering our findings of significantly higher privacy for one person, and the significantly lower privacy for two, compared to none (or one), we decided to code, posthoc (i.e., not as part of the experimental design) whether the people in the photo were looking at the camera or not. Our reasoning was that by looking at the camera, the subjects indicate they were aware that the photo was being taken, while people not looking at the camera may not be aware of the photo, which could affect perceptions of privacy about the image. Subjects were looking at the camera in about half of all the images with one person, and in about one-third of all the images with two people. We reran statistical models using the covariates as in the self-rated privacy model (#5) for each of the dependent variables, adding the variable “looking,” and found that it significantly reduced the likelihood of an image being perceived as private for the dependent variables Took Photo and In Photo/Own Home [coefficient = $-.55$ (95% CI: $-.82, -.09$); coefficient = $-.58$ (95% CI: $-.95, -.21$), respectively]. There was no statistically significant effect of looking in the model for Saw Photo [coefficient = $-.33$ (95% CI: $-.71, .06$)]. All other variables remained substantively and significantly the same as in the models reported in Tables 4–6.

This posthoc finding—that images of people looking at the camera are nearly 50% less likely to be perceived as private compared to images in which the people are not looking (all else equal)—is another indicator of socially embedded norms about privacy and appropriate image-sharing online. However, this finding can be considered only suggestive given that it was not part of the research design. The explosion of “selfies” posted online means that many more images will show people looking directly at the camera. However, as cameras become increasingly embedded in the environment, including being worn continuously by the people around us (e.g., through devices such as the new Google Clips camera, Snapchat Spectacles,¹² or the FrontRow FR Wearable Lifestyle Camera¹³), we should expect more and more images to circulate online in which the subjects may not be aware they are captured in the photo. Such practices have implications for privacy: both individual preferences about images we share and privacy’s broader social norms. Future research should more directly test whether the gaze of subject(s) in a photo, including looking at the camera or not, influences perceptions of privacy of the image.

6 INSIGHTS AND DISCUSSION OF FINDINGS

Overall, our findings across multiple analyses are consistent with theoretical expectations from the sociology of privacy and the contextual integrity of privacy. Specially, they indicate that privacy norms for personal photos posted online not only exist, but also vary according to differences in social context. The features of social context identified here include the number of people in the image, the presence of computer monitors, and the subjects’ role in relation to the photo, but not the type of household room. Our findings that expectations of privacy vary relative to the role of the subject support the theoretical claim that norms of privacy are based on the variation in acceptable levels of access to oneself and others [6]. In addition to the role of social norms, we show that individual privacy preferences also matter, but that it may depend on which measure of preferences is used. For example, the classic Westin typology used in many studies showed no association with expectations of privacy for personal images, but other measures did. The strongest and most consistent finding is our measure of self-rated privacy: Those who self-rated as private were significantly more likely to rate images as private across all models. In addition, participants under the condition “Own home/ In Photo” who scored higher on the Jensen Online Privacy mean scale [38], and the IUIPC-collection measure proposed by Malhotra et al. [50], were each significantly more likely to rate images under those conditions as private. Future work should continue to test and revise measures of privacy preferences to identify valid and reliable measures.

While some measures of individual privacy preferences matter, social privacy norms for online images are robust to differences in individual preferences. This is an important finding demonstrating the sociological nature of privacy in contrast to privacy as merely an individual preference. This study supports Nissenbaum’s theory of privacy as having contextual integrity because it shows that expectations of privacy not only vary with contextual features of an image but are also collectively shared. It also supports the sociological definition of privacy as access to self and others in which privacy norms govern expectations about appropriate access. Here, contexts that indicated greater access to the self were consistently shown to have higher expectations of privacy. Our findings are consistent with theories of privacy as socially and culturally dynamic rather than merely an individual preference. These findings also suggest that the so-called “privacy paradox,” in which individual privacy management behavior appears inconsistent with stated privacy concerns [24], may result from the lack of specification of the social context in which behavior occurs. In the face of an abstract question about sharing personal information

¹²<https://www.spectacles.com>.

¹³<https://www.frontrow.com>.

without social context, people may imagine social situations in which norms dictate more or less privacy, but this will not necessarily predict sharing behavior in any specific context. Put another way, it is likely that studies with statistical models that examine privacy behavior but include little or no information about the social context are likely underspecified (omitted variables) and so both theoretical conclusions and practical implications drawn from such studies are problematic. Not only are such studies problematic from a statistical point of view, but they incorrectly imply that preferences without reference to context or other constraints (e.g., budget) can explain behavior [22, 62]. However, the research reported here as well as other work shows that context and constraints do matter. The findings reported here for the privacy of personal photos highlight the need to consider not only the preferences of the photographer but also those of the image subject(s), as well as multiple social context dimensions, including image features like the number of people, and broader privacy norms, when considering privacy perceptions and behavior.

Our findings have implications for technology design, indicating that we need a design that supports and enables behavior consistent with both privacy preferences and social norms (see also [53, 74]). They also have implications for public policy, such as how laws should require social media companies to adequately enable and inform users about information practices, such as who has access to photos posted online and the extent of tagging content in personal images, to enable them to behave in ways consistent with their preferences and normative expectations about access (e.g., [52, 57, 78]).

Prior exploratory studies of photo privacy [32, 33, 42] found that the presence of computer monitors/screens [42], the type of room depicted in the image [32], and the number of people present [32], affected perceptions of privacy, and then extrapolated from those findings to propose privacy-preserving technology design. Though we did not find that privacy expectations differed by type of room, others have used room characteristics to automatically detect room characteristics for facilitating privacy controls [75]. Though we found computer monitors in photos lower expectations of privacy, Korayem et al. [42] used computer vision techniques to detect monitors in lifelogging images and block them without the need of users to manually flag such images. However, further work is needed to understand privacy expectations for computer screens in images. If it is the case that computer screens are simply so commonplace that they are not considered particularly private, then we would expect no effect on privacy at all. Furthermore, it may be the actual exposure of information via screens, and not the screen itself, that increases concerns about privacy in images.

Finally, though not an explicit focus of this study, subject demographic characteristics are included in all models (data not shown, findings reported in footnote 10), but none (gender, race, age, education) were statistically significant. Nor did we find any significant effects of frequency of technology use on expectations of privacy. Though this study is not designed to produce population estimates of privacy, it suggests that collective expectations of privacy (i.e., privacy norms) may be more powerful than individual characteristics. Future research is needed to examine population-level estimates and patterns of privacy across demographic groups.

Overall, our results show that people share common expectations about the privacy of online images, and that such privacy norms are socially contingent and multidimensional. Given these findings, image content detection algorithms alone are unlikely to be adequate for enabling privacy management. At the same time, these algorithms may exacerbate privacy concerns by disrupting normative expectations about appropriate access to personal images. Our findings also show that relying on any one dimension of accessibility—whether “type” of information, user privacy preferences, or even social context alone—is too limited to make sense of online behavior, since it depends on privacy norms as well as individual preferences and strategies for privacy management (see also [6, 14, 53, 65]). Instead, we must remember that privacy is contextual, multidimensional,

and socially contingent [6, 60, 68, 73, 91]. One of the socially contingent dimensions of privacy is normative agreement about appropriate access and information flows. Much as how legal scholars Warren and Brandeis worried that the emergence of instantaneous photographs in the early 20th century “invaded the sacred precincts of privacy and domestic life” [80], we consider the impact of widespread online sharing of personal images for privacy norms. Disrupting privacy norms about images of personal homes can have consequences for behavior and interaction, as well as the broader cultural meaning of privacy [34]. For example, as embedded and wearable cameras become more commonplace, expectations of privacy in public may decline, or bystander concerns may become so great that people begin to avoid certain places or people [63]. To the extent that places with embedded cameras are unavoidable (such as court houses, public transportation), people will not be able to act on their privacy concerns. More importantly, the consequences of such changes are not equitably distributed across population groups [6, 10].

7 LIMITATIONS

In evaluating the findings of this study, it is important to consider a number of scope conditions, as well as other limitations, that affect our results and their possible implications.

One scope condition was that, in seeking to use authentic personal images, this study used publicly available photos from online sources. Though having the advantage of being drawn from the real population of personal photos posted online (to particular social media sites), the fact that these images are already shared online suggests they are viewed, at least by the users who posted them, as inherently less “private.” This sample may be expected to have relatively low expectations of privacy, making it a conservative test of our hypotheses. However, we found that expectations of privacy are related to both image features and the subjects’ relation to the photo. Future studies may also want to sample both publicly available and private images to allow for greater potential variability in expectations of privacy. In addition, future research may wish to compare photos across social media sites, since there may be a variation in what is posted as well as what is considered private between different sites. However, we are also aware that individuals are not always aware of the ways in which their “public” online content is used for research purposes [19]. We acknowledge the ethical complexities of these questions, and suggest further studies on research using such “authentic” images would be valuable to understand not only general privacy norms but also those in regard specifically to online research.

Using real photos shared online also has the limitation of leaving us as researchers with less control over content, and thus somewhat less precision across photos within conditions. Future work may wish to test other types of photos, including staged photos, to permit more control over the characteristics being experimentally explored. However, we are not convinced that staged photos for image-vignettes would be better since what is gained in precision may be sacrificed in authenticity.

Another potential advantage to using staged photos would be more control over the characteristics of the people in the images, including apparent age, gender, and race/ethnicity. We included photos including people with the apparent gender of male and female, but we limited the variation in age and race/ethnicity characteristics. We did so to limit the scope conditions for this first-round experimental study, by excluding all photos with children and of people with apparent race/ethnicity that was not White/Caucasian. We recognize that identifying “apparent race/ethnicity” is problematic [29], and that this narrow set of images is a limitation of the study’s design, and conditions the scope of our findings (in that they apply to expectations of privacy for images with White/Caucasian people only). Given the role of inequality in privacy, in which those with fewer resources or a lower social status typically have less privacy [6], as well as specific research that shows the high levels of surveillance of people of color [10], we expect that norms of

privacy may also be different for different groups. Although we found no differences in perceptions based on the race/ethnicity of the participants in our study, we cannot say whether perceptions would vary based on the race/ethnicity of the photo subjects. Future studies should explore explicitly whether expectations for image privacy vary depending on the race/ethnicity, age, and other important sociodemographic characteristics of photo subjects, and how these factors intersect with other contextual features such as location or activity. Future work should also consider ways to test hypotheses like those proposed here using measures of behavior (e.g., avoiding social network contacts who post images considered too private) rather than simply expressed expectations.

Finally, this study was conducted via Amazon's mTurk platform. While steps were taken to recruit a broad sample of respondents, and control for respondent characteristics (including limiting to US respondents), this population has been shown to be different from the general population as a whole, in particular with regard to privacy preferences and sensitivity [40, 51, 81]. However, it is important to point out that our study is not intended to provide population estimates of privacy expectations but rather to examine patterns and relationships between hypothesized factors and conditions. mTurk is widely used for such studies [53, 56]; future work, however, should explore these questions on other platforms that include different populations.

8 CONCLUSION AND FUTURE WORK

Overall, our findings show that in addition to having individual preferences about privacy, people share common expectations, or social norms, about the privacy of personal images of the home shared online. Specifically, personal images that show one person only are two to three times more likely to be rated as private than photos with no people, or two people. However, images that show particular rooms like bedrooms, which may have previously been regarded as highly private, were no more or less likely to be rated as private than living rooms, dining rooms, or kitchens. This suggests that privacy expectations about the front-stage and back-stage areas of the home may be changing. We also found that privacy norms for personal images depend on a person's relation to the image, independent of the content of the image. That is, people are much more likely to say a photo is private, regardless of content, if they are in it. As new technologies continue to emerge and spread, including online sharing of personal photos, as well as increasing use of embedded cameras and face-recognition software, social norms regarding acceptable aspects of access or information flows about our personal homes, lives, behavior, and even likenesses [6, 60] are likely to be disrupted.

Online sharing of personal digital images is now ubiquitous in societies around the world, and certainly in the United States, the setting of this study. High-quality digital cameras are now a standard feature in mobile smartphones. Social media platforms such as Facebook, Snapchat, and Instagram (the latter used by 28% of adult Internet users in 2015, including 55% of users between the ages of 18 and 29) provide their users with new affordances and audiences for image-sharing [16]. Alongside camera-enabled smartphones, dedicated wearable cameras such as the Google Clips and Snap Spectacles, as well as lesser known products such as YoCam,¹⁴ and the iON SnapCam,¹⁵ cater to dedicated photographic self-trackers/"lifeloggers" by collecting large quantities of images automatically throughout the day, without users having to actively "take a photo." The explosion of online sharing of so many personal images, particularly images that reveal previously unobserved aspects of private lives at such a scale, can alter perceptions of privacy, including privacy norms—the commonly shared expectations about aspects of access to ourselves and others that are considered appropriate or not [6].

¹⁴<http://www.getyocam.com>.

¹⁵<https://usa.ioncamera.com/snapcam/>.

Technologies and practices that alter access to previously private places, behavior, and information can disrupt our expectations about privacy or what Nissenbaum [60, 61] calls the contextual integrity of privacy. Such disruptions have implications for individual behavior, social dynamics, and cultural values, as well for inequality, since privacy and its invasion are not equitably distributed across society. A better understanding of the social aspects of privacy, and how they are changing, will help guide both the design of technical tools to assist users, and the policies and institutions necessary to ensure that new technologies are consistent with our social norms and values about access and appropriate use.

REFERENCES

- [1] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (Jan. 2015), 509–514.
- [2] Alessandro Acquisti, Leslie K. John, and George Loewenstein. 2012. The impact of relative standards on the propensity to disclose. *Journal of Marketing Research* 49 (Apr. 2012), 160–174.
- [3] Shane Ahern, Dean Eckles, Nathaniel Good, Simon King, Mor Naaman, and Rahul Nair. 2007. Over-exposed?: Privacy patterns and considerations in online and mobile photo sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 357–366. DOI: <https://doi.org/10.1145/1240624.1240683>
- [4] Taslima Akter, Bryan Dosono, Tousif Ahmed, Apu Kapadia, and Bryan Semaan. 2020. “I am uncomfortable sharing what I can’t see”: Privacy concerns of the visually impaired with camera based assistive applications. In *29th USENIX Security Symposium*. USENIX Association, Boston, MA. Retrieved from <https://www.usenix.org/conference/usenixsecurity20/presentation/akter>.
- [5] Irwin Altman. 1977. Privacy regulation: Culturally universal or culturally specific. *Journal of Social Issues* 33, 3 (1977), 66–84.
- [6] Denise Anthony, Celeste Campos-Castillo, and Christine Horne. 2017. Toward a sociology of privacy. *Annual Review of Sociology* 43, 1 (Aug. 2017), 1–21.
- [7] Denise Anthony, Timothy Stablein, and Emily K. Carian. 2015. Big brother in the information age: Concerns about government information gathering over time. *IEEE Security & Privacy* 13, 4 (2015), 12–19.
- [8] France Belanger and Robert E. Crossler. 2011. Privacy in the digital age: A review of information privacy research in information systems. *MIS Quarterly* 35, 4 (Dec. 2011), 1017–1042.
- [9] danah boyd and Eszter Hargittai. 2010. Facebook privacy settings: Who cares? *First Monday* 15, 8 (Aug. 2010). DOI: <https://doi.org/10.5210/fm.v15i8.3086>
- [10] Simone Browne. 2015. *Dark Matters: On the Surveillance of Blackness*. Duke University Press, Durham, NC.
- [11] Celeste Campos-Castillo and Denise L. Anthony. 2014. The double-edged sword of electronic health records: Implications for patient disclosure. *Journal of the American Medical Informatics Association* 22, e1 (Jul. 2014), e130–e140.
- [12] Jaeyoung Choi, Martha Larson, Xinchao Li, Kevin Li, Gerald Friedland, and Alan Hanjalic. 2017. The geo-privacy bonus of popular photo enhancements. In *Proceedings of the 2017 ACM on International Conference on Multimedia Retrieval*. ACM, New York, NY, 84–92.
- [13] Robert B. Cialdini, Carl A. Kallgren, and Raymond R. Reno. 1991. *A Focus Theory of Normative Conduct: A Theoretical Refinement and Reevaluation of the Role of Norms in Human Behavior*. Elsevier.
- [14] Sunny Consolvo, Ian E. Smith, Tara Matthews, Anthony LaMarca, Jason Tabert, and Pauline Powledge. 2005. Location disclosure to social relations: Why, when, & what people want to share. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 81–90.
- [15] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In situ with bystanders of augmented reality glasses. In *The 32nd Annual ACM Conference*. ACM, New York, NY, 2377–2386.
- [16] Maeve Duggan. 2015. *Mobile Messaging and Social Media - 2015*. Pew Research Center Technical Report. Pew Research Center, Washington, DC
- [17] Amitai Etzioni. 1999. *The Limits of Privacy*. Basic Books, New York, NY.
- [18] Lujun Fang and Kristen LeFevre. 2010. Privacy wizards for social networking sites. In *Proceedings of the 19th International Conference on World Wide Web*. ACM, 351–360.
- [19] Casey Fiesler and Nicholas Proferes. 2018. “Participant” perceptions of Twitter research ethics. *Social Media + Society* 4, 1 (March 2018), 205630511876336-14.
- [20] Gary A. Fine. 2001. Enacting norms: Mushrooming and the culture of expectations and explanations. In *Social Norms*, Michael Hechter and Karl-Dieter Opp (Eds.). Russell Sage, New York, NY, 139–164.
- [21] Claude S. Fischer. 1992. *America Calling: A Social History of the Telephone to 1940*. University of California Press, Berkeley, CA.

- [22] Vaibhav Garg, Kevin Benton, and L. Jean Camp. 2014. The privacy paradox: A Facebook case study. In *Proceedings of the 42nd Research Conference on Communication, Information and Internet Policy*.
- [23] Erving Goffman. 1959. *The Presentation of Self in Everyday Life*. Doubleday, New York, NY.
- [24] Jens Grossklags and Alessandro Acquisti. 2007. When 25 cents is too much—an experiment on willingness-to-sell and willingness-to-protect personal information. In *Proceedings of the 6th Workshop on the Economics of Information Security*.
- [25] Vincent Toubiana, Vincent Verdot, Benoit Christophe, and Mathieu Boussard. 2012. Photo-TaPE: User privacy preferences in photo tagging. *WWW 2012* (2012), 617.
- [26] Rakibul Hasan, Eman Hassan, Yifang Li, Kelly Caine, David J. Crandall, Roberto Hoyle, and Apu Kapadia. 2018. Viewer experience of obscuring scene elements in photos to enhance privacy. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM.
- [27] Eman T. Hassan, Rakibul Hasan, Patrick Shaffer, David Crandall, and Apu Kapadia. 2017. Cartooning for enhanced privacy in lifelogging and streaming videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshop on Computer Vision Challenges and Opportunities for Privacy and Security*. IEEE, 29–38.
- [28] Michael Hechter and Karl-Dieter Opp (Eds.). 2001. *Social Norms*. Russell Sage, New York, NY.
- [29] Melissa R. Herman. 2010. Do you see what i am? *Social Psychology Quarterly* 73, 1 (Jan. 2010), 58–78.
- [30] Christine Horne. 2001. Sociological perspectives on social norms. In *Social Norms*, Michael Hechter and Karl-Dieter Opp (Eds.). Russell Sage, New York, NY, 3–34.
- [31] Christine Horne, Brice Darras, Elyse Bean, Anurag Srivastava, and Scott Frickel. 2015. Privacy, technology, and norms: The case of Smart Meters. *Social Science Research* 51, C (May 2015), 64–76.
- [32] Roberto Hoyle, Robert Templeman, Denise Anthony, David Crandall, and Apu Kapadia. 2015. Sensitive lifelogs: A Privacy Analysis of Photos from Wearable Cameras. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, 1645–1648.
- [33] Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014. Privacy behaviors of lifeloggers using wearable cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, New York, NY, 571–582.
- [34] Sarah Igo. 2018. *The Known Citizen: A History of Privacy in Modern America*. Harvard University Press, Cambridge, MA.
- [35] S. Jana, A. Narayanan, and V. Shmatikov. 2013. A scanner darkly: Protecting user privacy from perceptual applications. In *Proceedings of the 2013 IEEE Symposium on Security and Privacy (SP) Conference*. IEEE, 349–363.
- [36] Guillermina Jasso. 2006. Factorial survey methods for studying beliefs and judgments. *Sociological Methods & Research* 34, 3 (Jun. 2006), 334–423.
- [37] Guillermina Jasso and Karl-Dieter Opp. 1997. Probing the character of norms: A factorial survey analysis of the norms of political action. *American Sociological Review* 62, 6 (Dec. 1997), 947–964.
- [38] Carlos Jensen, Colin Potts, and Christian Jensen. 2005. Privacy practices of Internet users: Self-reports versus observed behavior. *International Journal of Human-Computer Studies* 63, 1–2 (Jul. 2005), 203–227.
- [39] Leslie K. John, Alessandro Acquisti, and George Loewenstein. 2011. Strangers on a plane: Context-dependent willingness to divulge sensitive information. *Journal of Consumer Research* 37, 5 (Feb. 2011), 858–873.
- [40] Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara B. Kiesler. 2014. Privacy attitudes of Mechanical Turk workers and the U.S. public. In *Proceedings of the 10th USENIX Conference on Usable Privacy and Security*. 37–49
- [41] Peter F. Klemperer, Yuan Liang, Michelle L. Mazurek, Manya Sleeper, Blase Ur, Lujo Bauer, Nitin Gupta Cranor, Lorrie Faith, and Michael K. Reiter. 2012. Tag, you can see it!: Using tags for access control in photo sharing. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 377.
- [42] Mohammed Korayem, Robert Templeman, Dennis Chen, David Crandall, and Apu Kapadia. 2016. Enhancing lifelogging privacy by detecting screens. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 4309–4314.
- [43] Jessica Lake. 2016. *The Face that Launched a Thousand Lawsuits: The American Women Who Forged a Right to Privacy*. Yale University Press, New Haven, CT.
- [44] Yifang Li, Wyatt Troutman, Bart P. Knijnenburg, and Kelly Caine. 2018. Human perceptions of sensitive content in photos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*.
- [45] Yifang Li, Nishant Vishwamitra, Bart P. Knijnenburg, Hongxin Hu, and Kelly Caine. 2017. Effectiveness and users' experience of obfuscation as a privacy-enhancing technology for sharing photos. In *Proceedings of the ACM on Human-Computer Interaction*, Vol. 1. ACM, 1–24.
- [46] David Lyon. 1994. *The Electronic Eye: The Rise of Surveillance Society*. University of Minnesota Press, Minneapolis, MN.
- [47] David Lyon. 2007. *Surveillance Studies: An Overview*. Polity Press, Cambridge.

- [48] Marry Madden. 2012. *Privacy Management on Social Media Sites*. Pew Research Center Technical Report. Pew Research Center, Washington, DC.
- [49] Charles Madge. 1950. Private and public spaces. *Human Relations* 3, 2 (Jun. 1950), 187–199.
- [50] Naresh K. Malhotra, Sung S. Kim, and James Agarwal. 2004. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research* 15, 4 (Dec. 2004), 336–355.
- [51] Jenny Marder and Mike Fritz. 2015. The Internet's hidden science factory. Retrieved from <https://www.pbs.org/newshour/science/inside-amazons-hidden-science-factory>.
- [52] Kirsten Martin and Katie Shilton. 2015. Why experience matters to privacy: How context-based experience moderates consumer privacy expectations for mobile applications. *Journal of the Association for Information Science and Technology* 67, 8 (May 2015), 1871–1882.
- [53] Kirsten E. Martin and Helen Nissenbaum. 2017. Measuring privacy: Using context to expose confounding variables. *The Columbia Science & Technology Law Review* 18, 1 (2017), 176. DOI : <https://doi.org/10.7916/stlr.v18i1.4015>
- [54] Gary T. Marx. 2016. *Windows into the Soul: Surveillance and Society in an Age of High Technology*. University of Chicago Press, Chicago, IL.
- [55] Barrington Moore Jr. 1984. *Privacy: Studies in Social and Cultural History*. M.E. Sharpe, Armonk, NY.
- [56] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman M. Sadeh. 2017. Privacy expectations and preferences in an IoT world. In *Proceedings of the Symposium on Usable Privacy and Security*.
- [57] Gina Neff and Laura Robinson. 2012. The social matrix of the emergent web: Governance, exchange, participation, & engagement. *Information, Communication & Society* 15, 4 (2012), 449–454.
- [58] Christena Nippert-Eng. 2010. *Islands of Privacy: Selective Concealment and Disclosure in Everyday Life*. University of Chicago Press, Chicago, IL.
- [59] Helen Nissenbaum. 2004. Will security enhance trust online, or supplant it? In *Trust and Distrust in Organizations: Dilemmas and Approaches*, Roderick M. Kramer and Karen S. Cook (Eds.). Russell Sage Foundation, New York, NY, 155–188.
- [60] Helen Nissenbaum. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford Law Books, Palo Alto, CA.
- [61] Helen Nissenbaum. 2018. Respecting context to protect privacy: Why meaning matters. *Science and Engineering Ethics* 24, 3 (2018), 831–852. DOI : [10.1007/s11948-015-9674-9](https://doi.org/10.1007/s11948-015-9674-9)
- [62] Greg Norcie and L. Jean Camp. 2015. The price of privacy: An examination of the economic costs of abstention from social network. In *Proceedings of the Amsterdam Privacy Conference*.
- [63] Alfredo J. Perez, Sherali Zeadally, and Scott Griffith. 2017. Bystanders' privacy. *IT Professional* 19, 3 (2017), 61–65.
- [64] Sandra Petronio. 2002. *Boundaries of Privacy: Dialectics of Disclosure*. SUNY Press, New York, NY.
- [65] Aarathi Prasad, Jacob Sorber, Timothy Stablein, Denise Anthony, and David Kotz. 2012. Understanding sharing preferences and behavior for mHealth devices. In *Proceedings of the 2012 ACM workshop on Privacy in the Electronic Society*, 117–128.
- [66] Blaine A. Price, Avelie Stuart, Gul Calikli, Ciaran McCormick, Vikram Mehta, Luke Hutton, Arosha K. Bandara, Mark Levine, and Bashar Nuseibeh. 2017. Logging you, logging me: A replicable study of privacy and sharing behaviour in groups of visual lifeloggers. In *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*. Article 22. DOI : <https://doi.org/10.1145/3090087>
- [67] Yasmeen Rashidi, Tousif Ahmed, Felicia Patel, Emily Fath, Apu Kapadia, Christena Nippert-Eng, and Norman Makoto Su. 2018. "You don't want to be the next meme": College students' workarounds to manage privacy in the era of pervasive photography. In *Proceedings of the 14th Symposium on Usable Privacy and Security*. USENIX Association, Baltimore, MD, 143–157. Retrieved from <https://www.usenix.org/conference/soups2018/presentation/rashidi>.
- [68] Priscilla M. Regan. 1995. *Legislating Privacy: Technology, Social Values, and Public Policy*. University of North Carolina Press, Raleigh, NC.
- [69] Reinhard Schunck. 2013. Within and between estimates in random-effects models: Advantages and drawbacks of correlated random effects and hybrid models. *Stata Journal* 13, 1 (2013), 65–76.
- [70] Reinhard Schunck and Francisco Perales. 2017. Within- and between-cluster effects in generalized linear mixed models: A discussion of approaches and the xthybrid command. *Stata Journal* 17, 1 (2017), 89–115.
- [71] Stuart Shapiro. 1998. Places and spaces: The historical interaction of technology, home, and privacy. *The Information Society* 14, 4 (1998), 275–284.
- [72] Aaron Smith. 2017. Record shares of Americans have smartphones, home broadband. Retrieved from <http://www.pewresearch.org/fact-tank/2017/01/12/evolution-of-technology/>.
- [73] Daniel J. Solove. 2008. *Understanding Privacy*. Harvard University Press, Cambridge, MA.
- [74] Luke Stark. 2016. The emotional context of information privacy. *The Information Society* 32, 1 (2016), 14–27.
- [75] Robert Templeman, Apu Kapadia, Roberto Hoyle, and David Crandall. 2014. Reactive security. In *Proceedings of The 2014 ACM International Joint Conference*. ACM, New York, NY, 1297–1306.

- [76] Edison Thomaz, Aman Parnami, Jonathan Bidwell, Irfan Essa, and Gregory D. Abowd. 2013. Technological approaches for addressing privacy concerns when recognizing eating behaviors with wearable cameras. In *Proceedings of The 2013 ACM International Joint Conference*. ACM, New York, NY, 739–10.
- [77] Matt Tierney, Ian Spiro, Christoph Bregler, and Lakshminarayanan Subramanian. 2013. Cryptagram. In *Proceedings of The 1st ACM Conference*. ACM, New York, NY, 75–88.
- [78] Joseph Turow, Michael Hennessy, and Nora Draper. 2015. The tradeoff fallacy: How marketers are misrepresenting American consumers and opening them up to exploitation. *SSRN Electronic Journal* (2015). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2820060.
- [79] Dave Vannette. 2017. Using Attention Checks in Your Surveys May Harm Data Quality. Retrieved from <https://www.qualtrics.com/blog/using-attention-checks-in-your-surveys-may-harm-data-quality/>.
- [80] Samuel D. Warren and Louis D. Brandeis. [1890] 1984. The right to privacy [The implicit made explicit]. In *Philosophical Dimensions of Privacy: An Anthology*, Ferdinand Schoeman (Ed.). Cambridge University Press, Cambridge, 75–103.
- [81] Jill Weinberg, Jeremy Freese, and David McElhattan. 2014. Comparing data characteristics and results of an online factorial survey between a population-based and a crowdsourcing-recruited sample. *Sociological Science* 1 (2014), 292–310. DOI: [10.15195/v1.a19](https://doi.org/10.15195/v1.a19)
- [82] Alan F. Westin. 1970. *Privacy and Freedom*. Atheneum, New York, NY.
- [83] Alan F. Westin. 2003. Social and political dimensions of privacy. *Journal of Social Issues* 59, 2 (2003), 431–453.
- [84] Allison Woodruff. 2014. Necessary, unpleasant, and disempowering—reputation management in the internet age. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 149–158.
- [85] Anna Wu and Xiaolong Zhang. 2011. Temporal sensitivity for location disclosure through mobile photo-sharing. In *Proceedings of the 1st international workshop on Mobile Location-Based Service*. ACM, New York, NY, 67–70.
- [86] Yaxing Yao, Justin Reed Basdeo, Oriana Rosata McDonough, and Yang Wang. 2019. Privacy perceptions and designs of bystanders in smart homes. In *Proceedings of the ACM on Human–Computer Interaction*, Vol. 3. ACM, Article 59. DOI: <https://doi.org/10.1145/3359161>
- [87] Jun Yu, Zhenzhong Kuang, Zhou Yu, Dan Lin, and Jianping Fan. 2018. Privacy setting recommendation for image sharing. In *Proceedings of the 2017 16th IEEE International Conference on Machine Learning and Applications*. IEEE, 726–730.
- [88] Jun Yu, Zhenzhong Kuang, Baopeng Zhang, Wei Zhang, Dan Lin, and Jianping Fan. 2018. Leveraging content sensitivity and user trustworthiness to recommend fine-grained privacy settings for social image sharing. *IEEE Transactions on Information Forensics and Security* 13, 5 (Jan. 2018), 1317–1332.
- [89] J. Yu, B. Zhang, Z. Kuang, and D. Lin. 2017. iPrivacy: Image privacy protection by identifying sensitive objects via deep multi-task learning. *IEEE Transactions on Information Forensics and Security* 12, 5 (2017), 1005–1016.
- [90] Shoshana Zuboff. 2019. *The Age of Surveillance Capitalism*. PublicAffairs/Hachette, New York, NY.
- [91] Elia Zureik, L. Linda Harling Stalker, Emily Smith, David Lyon, and Yolande E. Chan (Eds.). 2010. *Surveillance, Privacy, and the Globalization of Personal Information*. McGill-Queen’s University Press, Montreal, QC.

Received April 2018; revised November 2019; accepted December 2019