Chapter 6

Privacy Beyond the Individual Level

Jennifer Jiyoung Suh and Miriam J. Metzger

Abstract This chapter examines privacy as a multilevel concept. While current conceptualizations of privacy tend to focus on the individual level, technological advancements are making *group privacy* increasingly important to understand. This chapter offers a typology of both groups and group privacy to establish a framework for conceptualizing how privacy operates beyond the individual level. The chapter describes several contemporary practices that influence the privacy of multiple actors and considers the dynamics of multi-stakeholder privacy decision-making. Potential tensions that exist between the rights and preferences of individual group members or between individuals and the group as a whole are also examined. Finally, recommendations for tools and other mechanisms to support collaborative privacy management and group privacy protection are provided.

6.1 Introduction

Early privacy theorists conceptualized privacy in terms of control. Westin [1], for example, defined privacy as “the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others” (p. 5), and Altman [2] defined it as “selective control of access to the self or to one’s group” (p. 18). Despite these scholars’ acknowledgment of groups and institutions, the vast majority of privacy scholarship that has accumulated since has focused on the individual level [2, 3]. Indeed, most conceptualizations of privacy view it as a matter for *individuals* to manage by controlling others’ access to their personally identifying information. In line with this conceptualization, social, legal, and ethical paradigms that dominate discussions about privacy are also focused on individuals’ interests—for example, individual autonomy and personal freedom from surveillance [4]. Consequently, the tools,

J. J. Suh (✉) · M. J. Metzger

Department of Communication, University of California at Santa Barbara, Santa Barbara, CA, USA

e-mail: suh@ucsb.edu; metzger@comm.ucsb.edu

© The Author(s) 2022

B. P. Knijnenburg et al. (eds.), *Modern Socio-Technical Perspectives on Privacy*,
https://doi.org/10.1007/978-3-030-82786-1_6

91

laws, and policies currently in place to help people manage privacy—such as offering privacy settings to control one’s information in social media, ensuring anonymity to protect individual identity, or obtaining informed consent to collect and use personally identifiable information—rarely consider risks and threats that affect privacy beyond the individual level [5].

As a consequence of the emphasis on individual privacy, understanding privacy at the *group* level has received little attention in the research literature [2, 6]. However, the era of social media, big data, and data analytics poses new threats to privacy for groups and collectives, in addition to individuals [7, 8]. Advances in information and communication technologies over the last two decades have simultaneously increased opportunities for social sharing of information while diminishing control over that information (e.g., posting group photos in social media), often resulting in clashes between multiple stakeholders over protecting or revealing an image or piece of information. They have also spurred practices that not only acquire and collect individuals’ data but also aggregate data to identify trends in human behavior for modelling and making predictions about groups and collectives. These practices are often invisible to individuals. For example, geolocation information collected via GPS signals, cell towers, Wi-Fi connections, or Bluetooth sensors can be used to identify and/or predict mobility of migrant groups [4, 9] or, in the case of data from a fitness app used by soldiers, to identify the locations of secret military operations [10]. Further, data from individuals can be aggregated for purposes of predictive analytics and group profiling, such as likely academic or job performance (used by admissions officers or employers), health or financial status (used by medical insurers or loan officers), or criminality (used by law enforcement), which may result in discrimination against particular groups of people [11, 12].

These new privacy threats to groups as well as to individuals make it clear that individuals alone cannot manage their privacy effectively through merely controlling the flow of their own information. Recognizing the limitations of viewing privacy only at the individual level is an important starting point for expanding current views about privacy, as well as its protection. But what exactly is group privacy? Does it differ from individual privacy, and if so, how? This chapter begins with discussing current conceptualizations of both groups and group privacy to establish a framework for understanding the complex landscape of privacy at multiple levels. It then describes practices that influence the privacy of multiple actors, who may or may not realize they are a part of a group. Next, it considers the dynamics of multi-stakeholder privacy decision-making and potential tensions that exist between the rights and preferences of individual group members or between a member and the group as a whole. Finally, the chapter concludes with recommendations for tools and other efforts that support collaborative privacy management and group privacy protection.

6.2 Types of Groups and Types of Group Privacy

Conceptualizations of groups can be conceived to lie along two axes or dimensions: (1) how the group is constituted and (2) whether people are aware of the group's existence or their group membership status. In addition, group privacy can be conceived either in terms of the privacy of the group as a whole or in terms of the privacy of group members. These distinctions have important implications for social, legal, and technological mechanisms to protect privacy and thus are explicated below.

6.2.1 *Types of Groups: Self-Constituted Groups and Algorithmically Determined Groups*

In discussing group privacy, Taylor [9] explains that how people understand group privacy likely depends on what they mean by “the group,” and recent efforts to reconceptualize privacy illustrate that there are at least two different types of groups that must be distinguished: *self-constituted groups* and *algorithmically determined groups*. Most people are familiar with self-constituted groups, which refer to collectivities that are recognized as groups by their members or by outsiders (e.g., Girl Scouts, fan clubs for K-pop groups, Rotary Club, etc.). These groups tend to be stable over at least some period of time. Algorithmically determined groups, on the other hand, are often not self-constituted; rather they are identified by algorithms and typically are associated with group-level information that is obtained for some specific purpose, such as marketing (e.g., groups of people who buy natural hair care products and share demographic or geographic characteristics) or law enforcement (e.g., people who spend a lot of time at bars and night clubs in a specific neighborhood). Algorithmically constituted groups are also ad hoc and thus are usually less stable than self-constituted groups because group membership status may change with any tweak of the algorithm [13].

A further consideration for these two types of groups is the degree to which they are aware they are members of the group. In self-constituted groups, members usually know that the group exists and that they are members. As such, these groups are said to be “self-aware” [14]. In contrast, because algorithmically determined groups are discovered or “created” by data analytic technologies, group members are typically unaware of the group's existence. Such groups are increasingly prevalent and important because data analytic strategies, such as group profiling and data mining, are used across many sectors (businesses, education, health, government, military, etc.), so people can be part of these groups without being aware of the group itself or that they belong to them. See also [15] who use the terms “active” versus “passive” groups to differentiate groups that are self-aware from those that are not. An important implication of this for privacy is that group members are unable to protect themselves when they are not aware that algorithmic

group profiling has occurred and/or when they cannot detect their own membership in a group.

6.2.2 *Types of Group Privacy: “Their” Privacy and “Its” Privacy*

When thinking about privacy risks for a group, both “their” privacy and “its” privacy must be considered [4]. The difference between these two types lies in whether group privacy is constituted by concern for the “privacies” of individual group members (“their” privacy) or the privacy of a group as a whole (“its” privacy). In the first type, the privacy concern could be that revelations about a group would expose the identities of individual group members in a harmful way, for example, with members of a group of political dissidents. With the latter type, the privacy concern rests on the group’s very existence being discovered by a nonmember. Protecting one type of group privacy does not necessarily protect the other. For example, while anonymization of individual-level data may protect “their privacy” (i.e., the privacy of individual group members), it does not protect “its privacy” (i.e., the privacy of the group as a whole) from being detected by outsiders.

The case of Strava illustrates this point well. Strava is a popular fitness app that allows users to record and share their exercise routines via smartphone and fitness trackers. The data are collected from individual users anonymously, but they are then aggregated to produce heatmaps of popular exercise routes. One group of heavy users turned out to be US military personnel, and in 2017, it was discovered that the heatmaps based on data from this user group could reveal the locations and patrol routines of secret military bases overseas to anyone, including our enemies [10]. The heatmap used data that were anonymized and thus did not reveal personal information about any individual. But while aggregating anonymized individual data can protect group members’ identity, such data still have privacy implications for groups that are identified or profiled by the technology. The revelation of where a military group is located puts *both* the group as a whole, as well as individual members of that group, at risk. The lesson here is that protecting “their” group privacy through data anonymization does not necessarily protect “its” group privacy.

“Its” privacy and “their” privacy apply to both types of groups discussed in Sect. 6.2.1. Let’s take the case of a self-constituted and self-aware group of political dissidents. People may be concerned about “its” privacy if they are worried that the group’s existence may be discovered by the government, leading to the group’s dissolution. At the same time, people may also be concerned about “their” privacy if they are afraid that the identities of group members can be discovered by leaders of a repressive regime, resulting in group members’ imprisonment or worse. With an algorithmically determined group that is not self-aware, such as a collection of people who buy similar products and share demographic characteristics (e.g., women in San Francisco who like natural hair care products), people could be concerned about

“their” privacy if they are afraid that individuals who are categorized as members of the group would receive unwanted targeted advertisements. One could also be concerned about “its” privacy here too if they are uncomfortable with the idea of algorithms being used by marketers to group people into this, or any other, group to send targeted advertisements based on algorithmically predicted preferences that were not directly shared with the marketers.

6.2.3 Distinguishing Between Types and Levels of Privacy

These conceptualizations of privacy have raised debates about the degree to which individual privacy and group privacy, including both “its” and “their” types, are distinct. Two issues arise here. The first is whether “their” privacy amounts to anything more than individual privacy. Some argue that “their” privacy is the collection or sum of the privacies of the individual group members and thus is simply individual privacy. Others, however, maintain that the “their” type of group privacy is more properly conceptualized in a gestalt manner, as a property over and above the collection of the privacies of the individuals comprising the group, and thus is not the same as individual privacy [12]. Adopting this perspective, Belanger and Crossler [6] define group information privacy concern as “group members’ normalized view of information privacy concerns, which can be higher or lower than the individual members’ concerns taken as a whole” (p. 1031).

Second, adding to the complexity, Floridi [16, p. 90] explains how the notions of “its” group privacy and individual privacy may also intersect by giving rise to the notion of groups as individuals:

There are some kinds of rights that belong only to a group as a group, not to a group insofar as it is constituted by individual persons who enjoy those rights. In this case, it is important to understand that the group itself acts as an individual, to which a right is attributed.

While in most cases it may be easy to see the privacy of a group as a whole (“its” version of group privacy) and the individual privacy of group members as distinct (e.g., the right of a group not to be discovered by outsiders versus an individual’s right not to be identified as a member of a group), recent empirical studies have shown that it is not so easy for people to psychologically differentiate individual privacy from the “their” version of group privacy [17]. However, algorithmically determined groups that are not self-aware argue that “their” (group) privacy and individual privacy are separable, at least in theory, because while people are incapable of defending their individual privacy in such groups, laws can do so by recognizing and protecting the “their” type of group privacy (e.g., class action lawsuits). So even if differentiating these categories is impossible at a psychological level, they can be meaningfully differentiated at a legal level. These debates have important implications for the privacy rights of groups and individuals, which will be discussed in Sect. 6.5. In any case, perhaps the best that can be said is that while privacy is a multilevel concept, individuals are always important [6].

6.3 Contemporary Practices That Influence the Privacy of Multiple Actors or Groups

Technological advancements in recent years have enabled new practices that draw group privacy to the foreground. Such practices span across small group, organizational, and societal levels, affecting social groups, teams, as well as larger organizations and collectives. In all cases, these practices influence the privacy of multiple actors, rendering individual-level privacy insufficient to fully understand or to protect privacy in these contexts. Below are some examples that affect privacy at these various levels of analysis:

Example 1 One common and relatable example that illustrates how privacy risk can influence multiple actors in small groups is the practice of sharing group photos and tagging other users in social media posts. Most popular social network sites (SNSs) allow one user to share information about other users by posting group photos or by tagging them (e.g., User A tags User B, which then associates User B with User A's posts). As a result of these practices, one user has control over other people's information because there is not yet an effective tool or strategy that allows everyone involved to contribute equally to the decision of sharing a post about a group of users. In this situation, if the user who posts information about a group does not care about the privacy—and by extension the public image—of other members in the group (e.g., posting a group photo in which User A is shown in a positive light, but the others in the photo are not), this user will share the post and other members in the group will lose control over their information [3, 18]. Everyone in a group photo could have opinions about whether and how they want the photo to be shared online, but their personal opinions, or the group's collective opinion for that matter, are not taken into account.

Example 2 Workplace teams often use communication platforms that are administered by their organizations. While individual employees can use personal devices for private conversations, teams that discuss work-related matters that are not ready to be shared with the entire organization (e.g., special projects, secret assignments, etc.) often create private channels on communication platforms provided by their organizations (e.g., Slack, Microsoft Teams, etc.). While using these private channels for group conversations is efficient, individual team members have limited control over group-level information. For example, while a team's private channel may appear as "closed" to other employees, the fact that private channels are visually labeled as closed to others means that group itself (i.e., its existence) can be easily discovered. Indeed, and perhaps as a result of this, Microsoft recently announced that moving forward "private teams" on Microsoft Teams cannot be set as discoverable [19].

While this change may help work groups remain hidden to other employees, these communication platforms are administered by their organizations, so private groups can still be discovered and exposed to administrators that monitor and

regulate the use of these communication platforms across teams in the organization. One workaround is to use personal devices as an alternative private team channel. However, this is not enough to protect groups' privacy when employees are asked to install productivity monitoring software and/or security software to protect sensitive information about their organizations even on personal devices that are used to access work-related information (e.g., work email, etc.) [20, 21]. This example demonstrates how increasing workplace surveillance not only threatens personal privacy but also group privacy.

Example 3 The examples so far include people's decision to share information involving more than one person or a group of people, but individuals' actions to log data about their own personal behavior (e.g., lifelogging) can also affect group privacy. Lifelogging involves tracking personal data generated by our own behavioral activities. As more people use mobile and/or wearable devices, lifelogging has become very easy because many of these devices capture data about people's activities automatically (e.g., number of steps taken each day, details about workout routines or routes, etc.). Individuals' decision to log their lives may seem like it has nothing to do with group privacy, but the networked nature of individuals' data that are collected may expose lifeloggers to group privacy risks. A good example is Strava, the fitness app described earlier, which produced heatmaps that could compromise classified military information (e.g., strategic bases) and thus make groups and individuals, including military personnel and units on those bases, vulnerable to outside attacks. The Strava example is one of the few publicized cases, but it is not likely the only case because anonymized but aggregated location data are being used widely, and policy decisions based on such data could impact vulnerable groups that move with GPS-enabled devices, such as victims of natural disasters, patients fleeing from disease outbreaks, political asylum seekers and refugees, etc. [9].

Example 4 Lastly, emerging privacy threats from group inference technologies can even affect groups that have never shared anything about their group membership or information. Recent developments in AI-based group profiling and machine learning techniques enable marketers to not simply rely on data collected directly from their customers to design more effective targeted advertisements but by using technologies that make inferences about new potential consumer groups to target based on analyzing big data from a variety of sources. An example of these new techniques is a machine learning technology that correlates topics discussed on Twitter (e.g., #organicsampoo, #botanicalshampoo) with publicly available personal data of individuals who post about such topics (e.g., women who are in the ages of between 18 and 25 and live in San Francisco) [17]. This tool allows companies to make inferences about who and where they are likely to find new customers and thus to whom they should target their marketing messages. In other words, algorithmically determined groups could be used to draw inferences about potential customers for anyone who shares similar group (demographic and/or geographic) characteristics with the algorithmically discovered groups.

The important point to grasp here is that, based on this technology, groups of people who merely share particular demographic and geographic characteristics with other people who happen to discuss a topic in social media (e.g., natural hair care products) would receive targeted advertising or marketing messages about that topic. While sending targeted messages to people that share characteristics with a company's existing customer base may not seem like a new advertising or marketing strategy, the use of big data analytics makes the scale and reach of such messages both more invasive and pervasive than ever before. One seemingly benign hashtag that does not contain any personally identifiable information can, when aggregated, help companies extract group-level information that affects the lives of people who have never consented to sharing their group-level information with the data gatherer. While the information in this example is about hair products, the sensitivity of group-level information that is collected can vary, and the severity of negative consequences associated with different kinds of group-level information (e.g., identifiers for socially vulnerable groups, such as political protestors, sexual minorities, etc.) would vary accordingly.

The examples above illustrate why individuals cannot manage privacy by themselves, as well as when they should be concerned about privacy at the group level. And as technologies that enable group communication, creation, and discovery continue to advance, the number and types of practices that put the privacy of multiple stakeholders at risk will be even more far-reaching.

6.4 Dynamics of Multi-stakeholder Privacy Decision-Making

People are beginning to realize that they cannot effectively manage their own privacy by themselves because other group members' actions influence their own privacy, and in turn, they influence other group members' privacy. In other words, managing privacy is often not intrapersonal but *interpersonal* [22, 23]. The example of photo sharing in social network sites discussed in Sect. 6.3 illustrates how individual users do not have full control over their information because a group photo is co-owned. Indeed, participants in a recent survey reported that they preferred not to be tagged at all in photos because they want to be able to control their information [8]. The interpersonal nature of privacy in these kinds of scenarios raises the question of how to coordinate group members' expectations about appropriate information flow. Yet ways to collectively manage privacy with other people is not sufficiently addressed by how most privacy management options currently work, namely, individually managing control over one's own information through privacy settings.

A common problem that people experience from content generated by others that includes information about their own group belonging is face threats [24]. Face threats are verbal or nonverbal communication acts that challenge a person's self-presentation, and their consequences can vary in severity. For example, a post with multiple users tagged might reveal an individual's association with a social issue

that they had not been previously public about. Or, a group photo of teenagers at a party posted on social media might inadvertently reveal drug use. Research shows that people desire an effective way to manage privacy collaboratively in relation to face threats [25]. In fact, many people resort to relying on “mutual considerations” or “mental strategies” to do so. These strategies involve group members exerting mental effort during decision-making about whether to share information that they feel might cause face threats for other members and trusting that other members will do the same for them [26]. People often rely on these mental coping strategies to deal with privacy threats from others because they do not have alternative options.

However, while these mental strategies involve thinking about others’ desired self-presentation, relying on one person’s assumptions about what other group members would want is not always successful in reducing face threats because of misunderstanding, miscommunication, and mistaken assumptions [26]. Moreover, group members themselves may be concerned about whether they or others could actually succeed in living up to mutual expectations of making the right decision for each post and for everyone involved in a group [26].

Other studies show that people are starting to devise collaborative strategies to manage group privacy as co-owners of group-related information [27, 28]. Research by De Wolf et al. [28] suggests that members of groups may take the time to communicate, negotiate, and agree on what type of co-owned group information can be shared. For example, these researchers found that members of a youth organization in Flanders deal with group privacy management by employing a variety of communication strategies to coordinate privacy rules about their group, including *group privacy guidelines* (having explicit rules about what types of group information members can post on Facebook), *encryption* (interacting in a language that outsiders cannot understand), and *information management* (omitting information that one feels may anger other group members).

Cho and Filippova [27] aimed to create a comprehensive account of the types of privacy co-management strategies people use on Facebook. They found four strategies that people use to co-manage shared information. *Corrective strategies* include things like untagging or asking peers to remove content that allows users to control the visibility of content posted about them by others after the content has been published. *Preventive strategies* constrain the audience for shared information and may be enacted by using the friend lists feature to share content with a chosen group of people or by creating secret groups to share content. *Collaborative strategies* involve explicit coordination mechanisms to collectively manage each other’s privacy through negotiation. Similar to the members of the Flemish youth organization in [28], participants in Cho and Filippov’s study engaged in deliberate communication with each other about ways to manage their collective privacy. These included negotiating “rules of thumb” with their friends about sharing content concerning their group or discussing the appropriate privacy settings with their friends prior to disclosing content. Finally, *information control* is achieved by either self-censorship (as also found by [28]) or by making peace with the public nature of information sharing on social media. The most commonly applied privacy co-

management strategy was information control, which was followed by preventive, collaborative, and corrective strategies.

Jia and Xu [25] studied adoption of collaborative privacy management strategies by groups of linked contacts in social network sites. They found evidence for three types of rules negotiated by co-owners of shared information:

1. *Ownership management* rules that “define who the co-owners of the shared information are, with the assumption that co-owners should be able to make decisions about future disclosure of the collectively owned information” (p. 4289). This includes group members negotiating which group-owned information may be disclosed to others.
2. *Access management* rules that regulate disclosure and concealment of shared information to outsiders, ranging from open access to closed access. This may include coordinated content removal, restricting the visibility of shared information, or collectively deciding to provide unrestricted access to group information.
3. *Extension management* rules that govern decisions about whether to allow outsiders into the group privacy boundary by, for example, re-sharing group information by one member to people outside the group or adding new members to the group. Adopting and upholding these privacy co-management rules was positively related to a group’s collective value on privacy, the amount of disclosure of private information in the group, and group members’ perceived collective privacy risk.

Most of the research discussed so far takes the perspective that group norms shape rules that are developed within groups about whether and how to reveal or conceal collectively held information. This notion is central to the Theory of Multilevel Information Privacy (TMIP) proposed by Belanger and James [7] to understand how groups and individual group members make decisions about co-owned information. The theory posits that different social units (e.g., groups or individuals) can have different sets of rules about how to manage the unit’s information and interactions to protect privacy and also recognizes that people belong to multiple groups. Rule sets are thus activated according to the social identity that is salient in the decision moment. The social identity that is salient depends on the environment and specific situation or context. People will follow the normative rules of their social unit unless their privacy calculus (i.e., analysis of risks and benefits) indicates they should not. After a decision is made, positive and negative feedback shapes and refines their privacy rules and norms, which can affect future decisions. So, decisions about the same piece of co-owned information can be different in different environments, at different points in time, and if different social identities are made salient.

Engagement with the collaborative privacy management principles and strategies described in this section demonstrates that people are thinking about privacy at the group level and desire collaborative privacy management mechanisms. However, strategies that involve explicit group communication may not be enough to achieve effective group privacy management for many people because they are time-

consuming and may be uncomfortable to negotiate [27]. And some group members may not see the need for group privacy management, which can put other members at risk [28]. In general, people are more likely to use group privacy management strategies if they sense a stronger common bond as a group or feel highly attached to other group members [28]. Moreover, when people experience face threats as a result of others' privacy decisions, they often do not address the issue because they do not want to instigate conflict or "create drama," which they feel may hurt group cohesion [25, 29]. Perhaps for these reasons, collaborative group privacy strategies are still not as easily or widely applied as individual privacy management strategies that focus on individuals' controlling their own personal information through privacy settings.

Jia and Xu [25] moreover point out that many of the strategies for collaborative privacy management are only functional at a small scale or with a limited number of groups and become impractical and cost-inefficient in large social networks and when people interact with a large number of different groups. Another critique is that these strategies focus mostly on protecting individual privacy or the "their" type of group privacy rather than the "its" type of group privacy. For instance, healthcare teams using shared electronic medical records have rules to protect a patient's information (individual privacy), and friend groups in social media negotiate rules to avoid face threats of fellow members ("their" privacy). Although some collaborative privacy management strategies can be implemented to protect "its" group privacy, such as entirely closed access management rules or using preventative strategies to create secret groups (e.g., "finstas"), more typically the rules and strategies described in this section are used to make some particular *piece* of group information invisible to outsiders rather than to make the entire group itself undiscoverable. And perhaps most important, all of the strategies described above can only be used by groups that are self-aware. The privacy protection options for groups that are not self-aware, such as most algorithmically determined groups, are extremely limited. This issue will be addressed in Sect. 6.5.

6.5 Tensions Between Privacy Rights of Individuals Versus Groups

Whenever information is collectively held, tensions can arise about how to manage privacy. For example, individuals within groups may clash over their privacy preferences regarding information about the group. Take the example where one member wants to publish a group photo and another member does not, or where one wants the group as whole to be discoverable by outsiders, but another member does not. There is also the issue of privacy preferences of individual group members versus the group as a whole. Here, the group has negotiated and agreed upon a privacy rule (e.g., "no one shares information about our group to outsiders"), but then one group member violates the rule. Communication Privacy Management

theory [23] would discuss all of these examples in terms of privacy “boundary turbulence” among co-owners of information. Boundary turbulence is caused by a failure of privacy rule coordination between group members. It arouses negative emotions and has behavioral and relational consequences for co-owners [30].

Most of the research on boundary turbulence stays at the individual or at the “their” privacy levels, for example, by looking at how individual group members react to privacy breaches from other group members in terms of protecting their own (individual) or other group members’ (“their”) privacy. One example is when individual group member(s) withdraw from the member who caused the turbulence, through stonewalling, ignoring, or forcing the offending member out of the group [30]. Another example is if a group member withdraws their personal information from the group to protect their own privacy [16]. Turbulence also has group-level consequences, as it can prompt group members to collectively recalibrate, renegotiate, and re-coordinate their privacy rules [28, 31]. Much less research has examined the consequences of boundary turbulence in terms of individual group members’ relationships to groups as a whole (e.g., a member deciding to leave a group due to an instance of turbulence) or how boundary turbulence impacts the “its” type of group privacy. Boundary turbulence can threaten “its” group privacy (privacy of the group as a whole, such as its existence), with consequences that may be severe, including group infiltration, hostile takeover, harassment, or dissolution if all members decide to withdraw from the group.

Beyond preferences that may not align between individual members of a group or between a group and its members, there is also the question of the privacy *rights* of group members versus the group as a whole. While it is clear that group members have a right to privacy, this right is no different from individual privacy rights. More interestingly, there is debate about whether a *group* can have a right to privacy, or if that right is any different from the privacy interests of its individual members. Excellent discussions of this debate are available from [13, 32] and [33] (see also [17]). Bloustein [34] was the first to propose that groups have an interest in privacy. This interest, he argues, stems from group members’ desire to form associations privately with one another and legitimizes the idea of a group as a holder of privacy rights (rather than its members) because information about the existence of the group, and about the members who are associated with a group and with each other, can define a group’s identity in some cases.

Passive groups—groups that are not self-aware—complicate matters because if group members do not know that a group has been identified or that they are members, they have no ability to protect the group from unwanted intrusions. In these cases, how can or should privacy rights be protected for groups that are not self-aware? Some perspectives hold that a minimal level of “entitativity,” which is a perception of the extent to which a collection of people is perceived as a group by themselves or others, is a necessary condition for groups to have attitudinal and behavioral significance for people [35]. So clusters of individuals identified via an algorithm will not generate group privacy concerns if the individuals do not perceive themselves to constitute a group (see also [36]). In contrast, the minimal group paradigm in social psychology finds that mere categorization, even on an ad hoc

basis, reliably produces group identification and may lead to discrimination against group members [37]. This suggests that algorithmically determined groups who are not self-aware may (or should in theory) produce group privacy concerns for people and thus warrant claims that groups can be viewed as holders of privacy rights (see also [12]). Or in Mittelstadt's [33] words, "Algorithmically grouped individuals have a collective interest in the creation of information about the group, and actions taken on its behalf" (p. 475).

Finally, lawmakers have acknowledged the need to protect groups, even when members are not aware of their own group membership. Similar to algorithmically constituted groups, groups of people involved in class action lawsuits are ad hoc, and members may not have ever met or interacted with each other, but rather the group is constituted by a third party (i.e., the plaintiff) for a specific purpose (i.e., the lawsuit), and individuals' membership in the group is unbeknownst to them until they are notified of the lawsuit. Class action lawsuits are accepted as a critical tool to protect the interest of groups who do not have the means or ability to protect themselves from harms imposed on them by others. As such, they provide a legal framework for the protection of "its" group privacy rights even in the case of groups that are not self-aware [13].

6.6 Recommendations for Tools and Mechanisms to Protect Privacy Beyond the Individual Level

Protecting privacy beyond the individual is challenging because several parties are necessarily involved, which means communication, coordination, and, in some cases, conflict resolution are required. While some might assume that protecting individual privacy will protect group privacy, this is a fallacy. The Strava case is a good example of how group privacy can be compromised even when individual privacy is protected via anonymization. Kammourieh et al. [15] moreover argue that any privacy protection remedy based on individual identifiability is ineffective when the goal of an attacker is to identify or profile a group rather than to identify individuals. Identifying individuals is not necessary for group profiling. People may be acted upon in harmful ways through the act of being grouped, even without their personal identity being revealed and without knowledge that they have been categorized as a member of a group [9]. Because of this, groups need to safeguard their collective privacy and data protection rights [12]. To do so, new privacy protection solutions that are not exclusively based on individual privacy rights are needed [9, 13]. The remainder of this section offers ideas for some possible solutions.

- *Communication-based strategies for multi-party privacy management:* As discussed in Sect. 6.4, groups can engage in explicit deliberation to negotiate shared rules concerning how to protect the privacy of the group and/or group members. Examples include devising group privacy guidelines where group members

discuss whether and which content from a group event is appropriate to be shared on social media, encryption or using codewords and language that only group members know, and self-censorship [27, 28]. The downside to communication-based strategies to protect group privacy is that they are cumbersome and can be time-intensive, which likely explains why they are not widely used in practice. These strategies also do not apply to groups that are not self-aware.

- *Tools for multi-party privacy management:* While there are many tools available for people to protect personal privacy (e.g., privacy settings, anonymization and encryption of personal data, etc.), there are very few tools to protect group privacy. Yet such tools may help to overcome the overhead associated with adopting time-consuming collaborative group privacy management tactics, such as the communication-based strategies described above. Although still in their infancy, some prototypes exist for tools that allow multiple people to control content that involves more than one person [38–40]. For example, *CoPE* (Collaborative Privacy Management) is an application developed to aid collective privacy management of group photos on Facebook [40]. This tool alerts users to photos that they have been tagged in, requests and grants co-ownership of these photos, allows co-owners to see and change the privacy policies of individual pictures (i.e., control access to each photo), and provides photo browsing history. One drawback of the CoPE tool is that each co-owner separately specifies her or his own privacy preference for the shared photo instead of accommodating all stakeholders' privacy preferences or facilitating active negotiation of control between co-owners.

Another third-party Facebook application, *Retinue*, enables multiple associated users to specify their privacy concerns to co-control a shared group photo [38]. To resolve privacy conflicts caused by different privacy concerns of multiple users, a single data owner is specified who can take input from group members to make an appropriate privacy-sharing trade-off by adjusting the preference weights to balance the privacy risk and sharing loss for the group, taking all members' preferences into account. If a group member is not satisfied with the current level of privacy control, that user can adjust her/his privacy settings, ask the owner of the photo to change the weights for the privacy risk and the sharing loss, or report a privacy violation to request social network administrators to delete the photo.

Both CoPE and Retinue present usability problems for users, as they require extra layers of manual setting and re-setting of privacy preferences for shared content. A different approach to managing the privacy of co-owned information is “privacy nudges.” Nudges are short, on-screen, in situ messages that raise people's awareness of privacy issues. They are considered a “soft paternalistic approach” to increase user awareness of potential privacy risks and guide users to make more informed choices about their privacy management [41–43]. Although typically used to notify users about threats to their own privacy, nudges could be designed to help users become more aware about how their actions might impact group privacy. For example, a nudge might appear whenever users decide to tag

another person to confirm that they are indeed willing to share the co-owned information. A similar approach was proposed by [44], where users install a software daemon called *LocBorg*, which resides on a their computer or phone and protects the user from privacy violations by reminding them about the risks to their own and their groups' privacy in real time as they use social media apps such as Twitter.

- *Group privacy by design*: Group privacy management should not reside only in the hands of group members. Private companies could voluntarily develop their technology to be more accountable for protecting group privacy. "Group privacy by design" means designing products that incorporate protecting group privacy by default. Embedding group privacy management or protection tools into products (e.g., social network sites) is one form of this. An interesting idea to achieve group privacy by design is to use nudges to alert software engineers of potential dangers to group privacy during the design process to increase their awareness of vulnerabilities and, ideally, prompt them to eliminate dangers or insert group privacy protections as they develop systems and applications.

Group privacy by design is especially important for companies that use algorithms to identify groups from individuals' digital traces. Efforts to protect group privacy on the part of companies that use data analytic techniques to group people without their knowledge are needed because people who do not know they are being grouped cannot protect themselves from negative effects of such grouping. And companies stand to benefit from group privacy by design if it helps them avoid public outrage or boycotts from group privacy scandals. It is useful to recall that a good deal of the public outcry against Cambridge Analytica in 2018 was due to the company's failure to notify Facebook users that the company collected and processed not only individual users' data but also the data of users' linked contacts as well.

- *Self-regulation*: Companies should develop and then adhere to codes of ethical conduct to provide guidelines for responsible innovation, development, and usage of user data and algorithms to ensure the protection of both individual and group privacy. Rules surrounding the creation, accuracy, aggregation, deletion, storage, minimization, sharing, and other aspects of not just data collection but also its processing are important elements of such codes. Increasing transparency about classification and grouping algorithms, adopting policies that make clear to the public when and how group-related information is used, are also essential to effective self-regulation. Civil society groups such as consumer protection agencies and advocate organizations should be consulted during the development of ethical codes of conduct to help ensure privacy rights of groups are respected [45].
- *Government regulation*: Regulations governing data processing and the use of group-inference algorithms could prevent uses of data analytic techniques that profile and target groups and thus are another mechanism for protecting group privacy [15]. Specifically, policies that limit companies' use of sensitive

information pertaining to groups or require companies that use group-inference algorithms to provide notice to data subjects about how their information will be processed and/or to obtain informed consent from groups could be implemented. Policies about obtaining consent not just for data collection but also for any data processing or algorithms that may be applied to the data, such as those that aggregate anonymized datasets or use machine learning to infer group memberships, are useful. That said, obtaining informed consent from groups is difficult and could likely only be applied to groups that are, or are made to be, self-aware. Government mandates for companies to report their data processing methods alongside their potential risks to the public, as well as requiring procedures to allow users to opt out of data collection or aggregation, would go a long way toward group privacy protection. Allowing legal redress for violations to such policies is also important.

- *Education:* One of the major hurdles for protecting privacy beyond the individual level is the lack of public awareness about threats to group privacy [46]. Demystifying how big data analytics threaten not only individual but also the “their” and “its” types of group privacy would motivate people to begin to demand solutions. There are several means to educate the public on issues of data privacy, including campaigns by consumer protection agencies and advocates to increase awareness about the range of dangers algorithms pose to both individuals and groups; classes on data ethics, law, privacy, and digital rights in high schools and universities, especially early in data science training curricula; and continuing education for software developers [15]. Media reports of privacy scandals such as Strava and Cambridge Analytica also help raise public awareness of privacy threats posed to groups and their members.

6.7 Conclusion

Advances in technology in recent years have made privacy beyond the individual more pressing than ever. Group privacy has become increasingly important in the age of big data because most analytics target people not as individuals but rather as groups [13]. Groups, not individuals, are the object of value for data processors, as they care much less about a particular individual than they do about extracting behavior from individuals to shed light about groups who, for example, eat at different types of restaurants, prefer certain film or music genres, buy certain models or brands of cars, vote for liberal versus conservative candidates, are likely to suffer from a particular health issue, and so on. The privacy literature has been slow to recognize this, focusing instead on individual privacy interests, rights, and protection. A major purpose of this chapter has been to point out that by only protecting individual privacy, group privacy is not protected, and by revealing group privacy, individual privacy can be compromised. The implication of this co-

dependency is that *both* individual and group information must be protected in order to protect privacy effectively.

This chapter attempts to lay some of the groundwork for moving beyond the individual level in conceptualizing, theorizing about, and protecting privacy. By outlining how threats to privacy operate at multiple levels, providing examples of problems that people may experience as a result of threats to both “their” and “its” aspects of group privacy, and presenting recommendations for ways to resolve those problems, our hope is that this chapter will increase awareness and broaden the scope of scholarship on privacy that ultimately leads to more comprehensive and effective solutions to help both groups and individuals avoid privacy problems in the future.

References

1. Westin, A. 1967. *Privacy and Freedom*. New York: Atheneum.
2. Altman, I. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding*. Monterey, CA: Brooks/Cole Publishing Company.
3. Alsarkal, Y., N. Zhang, and H. Xu. 2018. Your privacy is your friend’s privacy: Examining interdependent information disclosure on online social networks. In *Proceedings of the 51st Hawaii International Conference on System Sciences*, pp 1–10.
4. Taylor, L., L. Floridi, and B. Van Der Sloot. 2017. Introduction: A new perspective on privacy. In *Group Privacy: New Challenges of Data Technologies*, ed. L. Taylor, L. Floridi, and B. van der Sloot, 1–12. Cham: Springer.
5. Cohen, J.E. 2012. *Configuring the Networked Self: Law, Code, and the Play of Everyday Practice*. New Haven, CT: Yale University Press.
6. Bélanger, F., and R.E. Crossler. 2011. Privacy in the digital age: A review of information privacy research in information systems. *Management Information Systems Quarterly*.
7. Bélanger, F., and T.L. James. 2020. A theory of multilevel information privacy management for the digital era. *Information Systems Research*. <https://doi.org/10.1287/isre.2019.0900>.
8. Birnholtz, J., M. Burke, and A. Steele. 2017. Untagging on social media: Who untags, what do they untag, and why? *Computers in Human Behavior* 69: 166–173. <https://doi.org/10.1016/j.chb.2016.12.008>.
9. Taylor, L. 2017. Safety in numbers? Group privacy and big data analytics in the developing world. In *Group Privacy: New Challenges of Data Technologies*, ed. L. Taylor, L. Floridi, and B. van der Sloot, 13–36. Cham: Springer.
10. Tufekci, Z. 2018. The latest data privacy debacle. *New York Times*.
11. Barocas, S., and H. Nissenbaum. 2014. Big data’s end run around anonymity and consent. In *Privacy, Big Data and the Public Good: Frameworks for Engagement*, ed. J. Lane, V. Stodden, S. Bender, and H. Nissenbaum, 44–75. New York, NY: Cambridge University Press.
12. Taylor, L., L. Floridi, and B. van der Sloot. 2017. *Group Privacy: New Challenges of Data Technologies*. Cham: Springer.
13. Floridi, L. 2017. Group privacy: A defence and an interpretation. In *Group Privacy*.
14. Taylor, L., B. van der Sloot, and L. Floridi. 2017. Conclusion: What do we know about group privacy? In *Group Privacy*.
15. Kammourieh, L., T. Baar, J. Berens, E. Letouzé, J. Manske, J. Palmer, D. Sangokoya, and P. Vinck. 2017. Group privacy in the age of big data. In *Group Privacy: New Challenges of Data Technologies*, ed. L. Taylor, L. Floridi, and B. van der Sloot, 37–66. Cham: Springer International Publishing.

16. Child, J.T., P.M. Haridakis, and S. Petronio. 2012. Blogging privacy rule orientations, privacy management, and content deletion practices: The variability of online privacy management activity at different stages of social media use. *Computers in Human Behavior* 28: 1859–1872.
17. Suh, J.J., M.J. Metzger, S.A. Reid, and A. El Abbadi. 2018. Distinguishing group privacy from personal privacy: The effect of group inference technologies on privacy perceptions and behaviors. *Proceedings of the ACM on Human Computer Interaction* 2: 1–22. <https://doi.org/10.1145/3274437>.
18. Yu, L., S.M. Motipalli, D. Lee, P. Liu, H. Xu, Q. Liu, J. Tan, and B. Luo. 2018. My friend leaks my privacy. In *Proceedings of the 23rd ACM Symposium on Access Control Models & Technologies (SACMAT)*, 93–104. New York, NY: ACM.
19. Microsoft. 2020. Manage discovery of private teams in Microsoft Teams. In *Microsoft Docs*. <https://docs.microsoft.com/en-us/microsoftteams/manage-discovery-of-private-teams>. Accessed 11 Aug 2020.
20. Chyi, N. 2020. The workplace-surveillance technology boom. *Slate*.
21. Roberts, J.J. 2020. Workplace privacy and surveillance software: What the law says | Fortune. *Fortune*.
22. Laufer, R.S., and M. Wolfe. 1977. Privacy as a social issue: A multidimensional development theory. *Journal of Social Issues*.
23. Petronio, S. 2002. *Boundaries of Privacy: Dialectics of Disclosure*. Albany, NY: State University of New York Press.
24. Litt, E., E. Spottswood, J. Birnholtz, J. Hancock, M.E. Smith, and L. Reynolds. 2014. Awkward encounters of an “other” kind: Collective self-presentation and face threat on facebook. In *Proceedings of the 17th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW’14)*. Baltimore, MD: ACM.
25. Jia, H., and H. Xu. 2016. Autonomous and interdependent: Collaborative privacy management on social network sites. In *Conference on Human Factors in Computing Systems – Proceedings*.
26. Lampinen, A., V. Lehtinen, A. Lehmuskallio, and S. Tamminen. 2011. We’re in it together: Interpersonal management of disclosure in social network services. In *Annual Conference on Human Factors in Computing Systems*, 3217–3226. <https://doi.org/10.1145/1978942.1979420>.
27. Cho, H., and A. Filippova. 2016. Networked privacy management in Facebook: A mixed-methods and multinational study. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW ’16)*, 503–514. San Francisco, CA: ACM.
28. De Wolf, R., K. Willaert, and J. Pierson. 2014. Managing privacy boundaries together: Exploring individual and group privacy management strategies in Facebook. *Computers in Human Behavior* 35: 444–454. <https://doi.org/10.1016/j.chb.2014.03.010>.
29. Wohn, D.Y., and E.L. Spottswood. 2016. Reactions to other-generated face threats on Facebook and their relational consequences. *Computers in Human Behavior* 57: 187–194. <https://doi.org/10.1016/j.chb.2015.12.021>.
30. Aloia, L.S. 2018. The emotional, behavioral, and cognitive experience of boundary turbulence. *Communication Studies*. <https://doi.org/10.1080/10510974.2018.1426617>.
31. Steuber, K.R., and R.M. McLaren. 2015. Privacy recalibration in personal relationships: Rule usage before and after an incident of privacy turbulence. *Communication Quarterly*. <https://doi.org/10.1080/01463373.2015.1039717>.
32. Floridi, L. 2014. Open data, data protection, and group privacy. *Philosophy and Technology* 27: 1–3. <https://doi.org/10.1007/s13347-014-0157-8>.
33. Mittelstadt, B. 2017. From Individual to group privacy in big data analytics. *Philosophy and Technology*. <https://doi.org/10.1007/s13347-017-0253-7>.
34. Bloustein, E.J. 1976. Group privacy: The right to huddle. *Rutgers-Camden Law Journal* 8: 219.
35. Campbell, D.T. 1960. Common fate, similarity and other indices of the status of aggregates of persons as social entities. In *Decisions, Values and Groups*.

36. Loi, M., and M. Christen. 2020. Two concepts of group privacy. *Philosophy and Technology*. <https://doi.org/10.1007/s13347-019-00351-0>.
37. Tajfel, H. 1970. Experiments in intergroup discrimination. *Scientific American*. <https://doi.org/10.1038/scientificamerican1170-96>.
38. Hu, H., Ahn, G.-J., and Jorgensen, J. 2011. Detecting and resolving privacy conflicts for collaborative data sharing in online social networks. In *Proceedings of the 27th Annual Computer Security Applications Conference – ACSAC '11*, 103. <https://doi.org/10.1145/2076732.2076747>
39. Squicciarini, A.C., M. Shehab, and F. Paci. 2009. Collective privacy management in social networks. In *Proceedings of the 18th International Conference on World Wide Web – WWW '09* 521. <https://doi.org/10.1145/1526709.1526780>.
40. Squicciarini, A.C., H. Xu, and X. Zhang. 2011. CoPE: Enabling collaborative privacy management in online social networks. *Journal of the American Society for Information Science and Technology* 63: 521–534. <https://doi.org/10.1002/asi.21473>.
41. Acquisti, A., I. Adjerid, R. Balebako, L. Brandimarte, L.F. Cranor, S. Komanduri, P.G. Leon, N. Sadeh, F. Schaub, M. Sleeper, Y. Wang, and S. Wilson. 2017. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys* 50. <https://doi.org/10.1145/3054926>.
42. Dogruel, L. 2019. Privacy nudges as policy interventions: Comparing US and German media users' evaluation of information privacy nudges. *Information, Communication & Society* 22: 1080–1095. <https://doi.org/10.1080/1369118X.2017.1403642>.
43. Solove, D.J. 2013. Introduction: Privacy self-management and the consent dilemma. *Harvard Law Review* 126: 1880–1903.
44. Zakhary, V., C. Sahin, T. Georgiou, and A.E. Abbadi. 2017. LocBorg: Hiding social media user location while maintaining online persona (vision paper). In *25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACM SIGSPATIAL 2017)*.
45. Zook, M., S. Barocas, D. Boyd, K. Crawford, E. Keller, S.P. Gangadharan, A. Goodman, R. Hollander, B.A. Koenig, J. Metcalf, A. Narayanan, A. Nelson, and F. Pasquale. 2017. Ten simple rules for responsible big data research. *PLoS Computational Biology*.
46. Metzger, M.J., J.J. Suh, S.A. Reid, and A. El Abbadi. 2021. What can fitness apps teach us about group privacy? In *Privacy Concerns Surrounding Personal Information Sharing on Health and Fitness Mobile Apps*, ed. D. Sen and R. Ahmed. IGI Global.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.