Ethics Inf Technol (2011) 13:289–302
DOI 10.1007/s10676-010-9224-8

ORIGINAL PAPER

Contextual gaps: privacy issues on Facebook

Gordon Hull · Heather Richter Lipford · Celine Latulipe

Published online: 30 April 2010

Abstract Social networking sites like Facebook are rapidly gaining in popularity. At the same time, they seem to present significant privacy issues for their users. We analyze two of Facebooks's more recent features, Applications and News Feed, from the perspective enabled by Helen Nissenbaum's treatment of privacy as "contextual integrity." Offline, privacy is mediated by highly granular social contexts. Online contexts, including social networking sites, lack much of this granularity. These contextual gaps are at the root of many of the sites' privacy issues. Applications, which nearly invisibly shares not just a users', but a user's friends' information with third parties, clearly violates standard norms of information flow. News Feed is a more complex case, because it involves not just questions of privacy, but also of program interface and of the meaning of "friendship" online. In both cases, many of the privacy issues on Facebook are primarily design issues, which could be ameliorated by an interface that made the flows of information more transparent to users.

Keywords Internet · Privacy · Contextual integrity · Social networking · Facebook

G. Hull (✉)
Department of Philosophy, University of North Carolina Charlotte, 9201 University City Blvd., Charlotte, NC 28223-0001, USA
e-mail: ghull@uncc.edu

H. R. Lipford · C. Latulipe
Department of Software Information Systems, University of North Carolina Charlotte, Charlotte, NC, USA
e-mail: Heather.Lipford@uncc.edu

C. Latulipe
e-mail: clatulip@uncc.edu

Introduction

A substantial and growing literature confronts privacy issues raised by the development and general dissemination of new information technologies. As Helen Nissenbaum puts it, privacy "has been a fixture in public discourse through radical transformations of technology from stand-alone computers … to the current distributed network of computers with linked information systems" (2004, 101). In what follows, we use Nissenbaum's analysis of privacy in terms of "contextual integrity" as a framework for understanding the privacy implications of recent developments on the social networking site (SNS) Facebook. Assessment of these developments is important. Apart from the growth of social networking more generally, well upwards of 80% of college students maintain Facebook profiles, and the expansion of the site beyond college campuses is reflected in its continued, rapid growth (over 400 million users as of this writing). One recent study found that Facebook use jumped 700% from April 2008 to April 2009, with users logging a total of almost 14 billion minutes on the service during the month of April 2009 (Nielsen 2009). As of March, 2010, Facebook is the second ranked site on the Internet traffic metrics on alexa.com, accounting for almost 5 percent of all global pageviews. Further, as we will detail below, there is considerable discussion of what the widespread adoption of SNS means for users' privacy. Facebook itself does a number of different things with users' data that suggest privacy issues.[1] We will discuss two: the sharing of users' friends' data

[1] For a discussion of Public Search, Social Ads, and Beacon, with reference to the company's terms of service, see Hashemi (2009). For an accessible overview of SNS, focusing on Myspace, see boyd (2007).

with third parties via "applications," and its switch to "News Feeds" (rather than page visits) as the primary way to get updates on one's friends. In both cases, Facebook's interface design, which tends to occlude the flow of information from one context to another, emerges as a significant contributor to privacy problems. That said, we will argue that the issues posed by Applications are considerably more difficult to resolve than those of News Feed.

Our purpose is accordingly "disclosive" in that our effort is toward "disclosing and evaluating embedded normativity in computer systems, applications, and practices" (Brey 2000, 12). Of course, that there are privacy issues in Facebook is not news. Our point with regard to disclosive ethics is twofold. First, the privacy implications are the result of design decisions in Facebook, even when those design decisions are not themselves obviously about privacy. Our choice of News Feed and Applications is meant to be exemplary of the sorts of normativity that can be embedded in social networking sites. Second, Helen Nissenbaum's account of privacy as "contextual integrity" offers a useful conceptual apparatus for understanding why and how these design decisions have privacy implications. In this regard, the point will not be so much to apply Nissenbaum's theory, as to use her theoretical vocabulary to make salient some of the reasons why News Feed and Applications generate the privacy complaints they do.[2]

The paper proceeds as follows. In the next section, "Privacy as contextual integrity", we outline Nissenbaum's contextual integrity framework in the context of social networking. In the following section, "How social networking sites change the privacy equation", we discuss social networking as a more general phenomenon, and two phenomena in particular—blogging and webcamming—that can be seen as precursors to the issues we discuss in Facebook. The following two sections then discuss Facebook's Applications and News Feed features. Throughout, we will attempt both to clarify what the privacy implications of those features might be, and to situate them partly as questions of interface and interaction design. We then offer some concluding thoughts.

[2] Nissenbaum says that her paper is not "aiming for a full theory of privacy, but only a theoretical account of a right to privacy as it applies to information about people," as that information is gathered in public surveillance (2004, 106). Our goal here is similarly narrow, and we remain agnostic about topics like the "essence of privacy." Solove also cautions against the (perhaps quixiotic) attempt at producing a full theory: "the quest for a traditional definition of *privacy* has led to a rather fruitless and unresolved debate. In the meantime, there are real problems that must be addressed, but they are either conflated or ignored because they do not fit into various prefabricated conceptions of privacy …. In this way, conceptions of privacy can prevent the examination of problems" (2007b, p. 759). For a survey of definitional attempts, see Nissenbaum (2004, pp. 107–113).

Privacy as contextual integrity

Helen Nissenbaum's work offers an analytical framework for understanding the commonplace notion that privacy is "contextual." Nissenbaum's account of privacy and information technology is based on what she takes to be two non-controversial facts. First, there are no areas of life not governed by context-specific norms of information flow. For example, it is appropriate to tell one's doctor all about one's mother's medical history, but it is most likely inappropriate to share that information with casual passers-by. Second, people move into and out of a plurality of distinct contexts every day. Thus, as we travel from family to business to leisure, we will be traveling into and out of different norms for information sharing. As we move between spheres, we have to alter our behaviors to correspond with the norms of those spheres, and though there will always be risks that information appropriately shared in one context becomes inappropriately shared in a context with different norms, people are generally quite good at navigating this highly nuanced territory. On the basis of these facts, Nissenbaum suggests that the various norms are of two fundamental types. The first, which she calls norms of "appropriateness," deal with "the type or nature of information about various individuals that, within a given context, is allowable, expected, or even demanded to be revealed" (2004, 120). In her examples, it is appropriate to share medical information with doctors, but not appropriate to share religious affiliation with employers, except in very limited circumstances. The second set are norms of "distribution," and cover the "movement, or transfer of information from one party to another or others" (2004, 122). She grounds these norms partly in work by Michael Walzer, according to which "complex equality, the mark of justice, is achieved when social goods are distributed according to different standards of distribution in different spheres and the spheres are relatively autonomous" (2004, 123). Information is such a social good. Thus, in friendship, confidentiality is a default rule: it may be appropriate to share the details of one's sex life with a friend, but it is not appropriate for that friend to broadcast the same information on her radio show. In medical situations, on the other hand, a certain sharing of information—from a radiologist to a primary care physician, for example—is both normal and expected.

Nissenbaum emphasizes two general points that emerge. On the one hand, information is always tagged, as it were, with the context in which it is revealed: there is no such thing as context-free information. On the other hand, the scope of privacy norms is always internal to a context. There is no such thing as a universal privacy norm. These results have striking theoretical implications. Most discussions of privacy, she argues, turn on one or more of

three basic principles (2004, 107–12). The first seeks to protect the privacy of individuals against intrusive government agents. The second seeks to limit access to information deemed sensitive or intimate. The third focuses on curtailing intrusion into spaces that are deemed private or personal. The problem is that none of these principles is adequate to describe the privacy concerns of public surveillance: such surveillance is often not conducted by government agents (instead being conducted by corporations or other private actors); it is often not about sensitive information; and it is generally not in private spaces. As a result, issues of public surveillance often slip under the radar of privacy theory based on these principles, appearing either not to be a problem or to be intractable (2004, 116). Nissenbaum's suggestion is that the limitations of the three principles derive from their effort to speak universally to all privacy situations, and from their tendency to treat information dichotomously as either private or non-private, with no middle ground (2004, 118). Nissenbaum's reinterpretation of the notions of information and its flow undergirding any privacy theory allows her to then develop an understanding of privacy that responds to the limitations she identifies with the three principles.

Nissenbaum's initial exploration of the contextual integrity framework confines the analysis to issues of public surveillance. Her examples are accordingly about government records (like court records) being placed online, consumer profiling and data mining, and the growing use of RFID tags. Each of these practices poses intuitive privacy issues that are difficult to capture with the three principles. For example, court records are by definition already public, and placing them online seems to be a case of government doing what it has always done, but more efficiently. Since governmental efficiency is presumptively good, and since the information is already public, it is hard to explain privacy discomfort over the practice (2004, 104). Nissenbaum proposes that the issue lies with norms of distribution: although public records have always been available at the local courthouse, the difficulty of traveling to the courthouse to get them tended to limit their exposure to those with a significant interest in them. Placing records online makes them readily available to those with no connection to the information and no particular interest in it.

In our view, the use of Nissenbaum's account extends well beyond the public surveillance context to which she applies it. Evidence from social networking sites, and Facebook in particular, suggests that contextual integrity will be an apt analytical framework in that context as well. First, and unlike popular perceptions and perhaps unlike other SNS, almost all of the evidence suggests that Facebook users primarily use the site to solidify and develop their offline social relationships, rather than to make new relationships online. Information on Facebook would predictably tend to mirror the offline tendency to contextual situatedness.[3] Second, evidence about social conflicts in Facebook suggests that one of the greatest sources of tension is when information that would remain internal to one context offline flows to other contexts online. This is in large part an interface issue—there is no easy way for a user to establish and maintain the many separate and sometimes overlapping social spheres that characterize offline life. If these fluid spheres are difficult to explicitly create on Facebook, then managing information flows among them is likely to be very difficult. Indeed, the underlying architecture of Facebook is largely insensitive to the granularity of offline social contexts and the ways norms differ between them. Instead, the program assumes that users want to project a single unified image to the world. These contextual gaps are endemic to SNS, and a substantial source of privacy problems.[4] For example, the 'status line,' generally gets projected indiscriminately to all friends. One can think of the status line as an answer to the questions "How's it going?" or "What's up?" However, in offline life one tailors responses to the audience one is with. While Facebook now allows users to create multiple friend groups and then control which groups can see their status, this requires explicit, cumbersome effort and falls far short of reflecting the many complex and overlapping spheres of offline life.

How social networking sites change the privacy equation

Two recent Web phenomena, which pre-date social networking sites, highlight some of the privacy issues that are

[3] See Joinson (2008), Lampe et al. (2008, 2007, 2006) and Ellison et al. (2007). Not all social networking sites will be the same; researchers at IBM noted that individuals tended to use an intra-corporate site primarily for making new contacts, and not for maintaining old ones (DiMicco et al. 2008).

[4] For a similar discussion of problems of "social convergence" on Facebook, see boyd (2008, 18–19). For the efforts of teenagers to deal with this problem on Myspace (in particular, the porous boundaries between "teen" and "parent" contexts), see boyd (2007). See also Binder et al. (2009) (suggesting that social networking sites tend to flatten out the separateness of offline social spheres, and that "privacy" concerns can best be understood as one part of this problem of tension between poorly demarcated social spheres). The flattening of contexts discussed here does not extend as far as the vertiginous "context collapse" confronting those posting YouTube vlogs, the audience for which is anyone, everyone, and no one all at once (Wesch forthcoming). This is so for two reasons: (a) the problem for vloggers is the inability to imagine a context. Facebook users do imagine a context for their postings, even if they get that wrong. (b) Wesch suggests that YouTube vloggers are not generally supporting offline social networks. As noted above, this is in sharp contrast with the evidence about Facebook.

relevant to Facebook and other social networking sites. The first is blogging. The number of people who post their thoughts online in Web logs is staggering, and growing rapidly. Technorati reported having indexed 133 million blog records in August, 2008; during one 24-h period, the site measured 900,000 individual blog posts.[5] Bloggers talk about all sorts of things; many, many of them use their blogs essentially as diaries. Since blog entries are both accessible and more or less permanent, they have significant implications for privacy—not just that of their authors, but also of those about whom they write.[6] One particularly striking example is the case of Jessica Cutler, alias Washingtonienne, who blogged in graphic detail about her affair with an attorney working for a US Senator.[7] She had a relatively small readership until her blog was picked up by the popular Washington gossip blog Wonkette. Although she was subsequently fired, Cutler enjoyed a newfound celebrity, including lucrative book deals; at one point she declared that "public embarrassment is really very liberating. You really stop caring about what people think, which is something only the elderly seem to able to accomplish with great aplomb" (Cutler 2004). The lawyer did not agree, and he filed multiple lawsuits against Cutler for invasion of privacy (as of this writing, the various legal actions surrounding the suit are still not fully settled). Of interest here is not so much to belabor the obvious point about the public nature of the Internet and norms of distribution; it is to note by hyperbolic example how easy it is for one's social networking to cause collateral damage (in slightly more delicate terms, a "spillover"). In other words, one of the reasons that norms of distribution matter is that adherence to them can protect third parties. Indeed, Nissenbaum lists among the values supported by these norms the prevention of informational harms, the preservation of human relationships, and the preservation of the space needed for individuals to develop autonomous life plans (2004, pp. 129–133).

Cutler's easy dismissal of embarrassment allows one to note a second, related phenomenon: webcamming, particularly by women. The most prominent example here is "Jennicam," in which Jennifer Ringley put webcams throughout her apartment and left them on all the time, broadcasting the quotidian details of her life to site visitors, who could see anything from an empty room, to Jenny typing at her computer, to her having sex with her boyfriend. The site spawned a considerable debate about the implications of such a practice for feminism and gender relations: was Jennicam's deliberate, nonchalant placement of cameras taking agency back from the male gaze, or pandering to it, reinforcing the patriarchal norm that women are fundamentally always on display for male consumption? We would argue that Jennicam provoked a needed conversation by making explicit the social norm according to which information flows from women's bodies to those silently observing them.[8]

Generally speaking, one might suggest that constant publicity tends to cause people to modify their desires and behaviors accordingly. This is because one's adaptation to social norms of information flow are important to fashioning the social personae through which one manages individual interactions (Nissenbaum 2004, 130). To the extent that such a phenomenon is happening, it has profound implications for personhood and basic ethical norms of social interaction.[9] In this context, we would like to emphasize two narrower points. First, whatever else webcamming does, it made the pervasive information flows apparent. There might have been problems with Jennicam, but one of them was not that Ringley had no idea that she was being watched. The flows of information from Ringley's life to the Internet were completely transparent. Second, norms are not static; in particular, they can be subject to feedback loops, where changes in the amount of privacy that one experiences can lead to changes not just in when one does or does not expect privacy, but how much privacy one wants. In other words, the urge to seek privacy can be changed by publicity into an urge to seek publicity; in an online environment, we can expect that individuals

[5] The exact numbers are difficult to quantify. One earlier study estimated that half of blogs are by children and teenagers (cites in Solove 2007a, 24). Technorati (2008), on the other hand, reports that 75% of bloggers have college degrees, and 42% have been to at least some graduate school.

[6] As Solove puts it, "as people chronicle the minutiae of their daily lives from childhood onward in blog entries, online conversations, photographs, and videos, they are forever altering their futures—and those of their friends, relatives, and others" (2007a, 24).

[7] The details of the story are taken from Solove (2007a), 50-4. As Solove puts it later, "even if information is already circulating orally as gossip among a few people, putting it online should still be understood as a violation of privacy—even if it is read only by people within one's social circle…. Putting the information online increases dramatically the risk of exposure beyond one's social circle. Placing information on the Internet is not just an extension of water cooler gossip; it is a profoundly different kind of exposure, one that transforms gossip into a widespread and permanent stain on people's reputations" (2007a, 181).

[8] For a summary of the debate that registers the substantial difficulties that Jennicam poses from a feminist standpoint, see Bailey (2009). For a celebration, see Jimroglu (1999).

[9] For personhood and identity formation in the context of SNS, see Papacharissi (2009) and Livingstone (2008); more generally, see, for example, Matthews (2008) and the work of Andy Clark, most accessibly in Clark (2003). For an initial attempt at discussing some of the ethical implications of Clark's analysis, see Hull (2009) (in the case of library filtering programs); for a Kantian application of these issues see Myskja (2008); for a thoughtful application to privacy, see Cocking (2008).

will come to present themselves differently, in order to take advantage of the inevitability of publicity.

These examples show privacy issues that emerge because the Internet is basically a global network in which everyone can participate. Information posted to the Internet is potentially visible to all. For most people, such universal broadcast of information has no parallel offline. In other words, offline personal information is seldom communicated to a context anywhere near as broad as the entire Internet. As social networking sites continue to grow and increasingly integrate with the rest of the Internet, they should be expected to inherit some of these issues. At the same time, the contextual situation for social networking is made both more complex and more difficult by the import of the offline concept of "friend." Information flows on social networking sites are mediated not just by the global nature of Internet communication, but by the ways that those sites and their users interpret the meaning of online friendship and the social norms that go with it. Not surprisingly, empirical research indicates that SNS users (possibly excepting younger teens on Myspace) tend to construct their identities relationally, becoming and expressing who they are by way of highlighting whom they associate with (Papacharissi 2009; Livingstone 2008).

While online friendship may be derived from offline notions of social relationships, "friends" online are clearly different from friends offline. As noted above, Facebook communities tend to have a basis in offline communities, and users appear largely to use the site to strengthen and reinforce those offline social ties. However, it is quite easy to "friend" someone, and Facebook users often have vast numbers of casual friends (the average user has 130 friends on the site, according to Facebook), far in excess of what they would be able to maintain offline. In the online context, Yochai Benkler suggests that two general phenomena can be observed. First, "we see a thickening of preexisting relations with friends, family, and neighbors, particularly with those who were not easily reachable in the pre-Internet-mediated environment" (2006, 356). Within these relationships, there is a general emphasis on peer relationships over hierarchy and familial ties. Second, "we are beginning to see the emergence of greater scope for limited-purpose, loose relationships" (ibid.), as for example those surrounding topic-specific blogs.[10] Facebook would seem to be a natural outgrowth of these developments; it appears to operate largely as an example of the first phenomenon, but the second phenomenon shows up in the form of groups and fan pages.

People join social networking sites for a variety of reasons. As noted above, empirical research suggests that most Facebook users utilize the site to keep up with offline friendships, or to reconnect with individuals from their past. In so doing, they develop relational aspects of their identities. More theoretically, one might say that SNS allow users to fill in the "fuzzy edges" of their social networks—those who have interests like their own, but whom they do not already know (Gelman 2009). By signaling not only that I am a graduate of Behemoth High, but that this fact about me is one that I consider relevant, I enable others with the same attributes to find me. Conversely, I can find others with interests similar to my own. Sites thus facilitate the formation of networks of people with similar interests, and thereby provide tremendous value to their users. Empirical research accordingly has found that, although the causality is difficult to determine, intensity of Facebook use is positively correlated with all kinds of social capital (Ellison et al. 2007).

Of course, in order for the sites to perform this function, individuals have to disclose information about themselves. Not surprisingly, therefore, Facebook and other SNS present well-publicized privacy concerns. It is now standard fare for college advising departments to remind students that their future employers are likely looking at their Facebook history; pictures of drunken partying have a way of causing job interviews not to happen. Despite these warnings and a number of well-publicized career-damaging incidents, Facebook users continue to post vast amounts of information about themselves. Gross and Acquisiti summarized the results of one of the first studies on the topic:

> It would appear that the population of Facebook users we have studied is, by and large, quite oblivious, unconcerned, or just pragmatic about their personal privacy. Personal data is generously provided and limiting privacy preferences are sparingly used. Due to the variety and richness of personal information disclosed in Facebook profiles, their visibility, their public linkages to the members' real identities, and the scope of the network, users may put themselves at risk for a variety of attacks on their physical and online persona (2005, 8).

[10] Benkler's analysis is much more complex than can be detailed here, and adduces relevant empirical support. Two claims that we will not address further are (1) "Facebook friends aren't real friends." As we argue below, the norms of online friendship are precisely what is at stake for sites like Facebook. Also, and as Benkler emphasizes, offline friendships are heavily conditioned by being in a highly fragmented, mass-media society. The Internet is not replacing utopia, and should not be compared to utopia in order to be found wanting. (2) "People waste time on Facebook." This complaint loses some of its force in the face of empirical data that suggests people take their Internet time from TV. Facebook friends may or may not be a person's "real" friends, but they're certainly a step closer to users than Ross and Rachel. According to Technorati (2008), the average blogger watches only a third as much TV as the average adult.

Users risk becoming victims of stalking, re-identification in other contexts (as, for example, by future employers), and even identity theft. Despite the risks, this behavioral finding appears to be quite robust. The nature of the problem seems to have changed over time: earlier research reported that users did not care about privacy; more recent research indicates that users care about privacy but, for various reasons, are unable to take effective action to protect it. Many of these reasons have to do with the tendency of Facebook to encourage mistakes in risk perception.[11] Additionally, since the pictures (for example) that one user posts often include images of her friends, the privacy problem is not just what the individual user does to herself, but what she does to her friends. Contextual integrity makes it immediately clear where the core problems lie. On the one hand, norms of appropriateness suggest that not all friends should be treated to the same amount of personal information; on the other hand, norms of distribution indicate that it is acceptable to share photos of one's social life with one's friends but that it is not acceptable to share those details with one's friends' employers.

Negotiating such privacy issues is a normal part of everyday life. Social networking sites, however, make such privacy negotiations both more salient and more troublesome. There are at least four reasons why this is so. First, as Strater and Lipford point out, "while managing identity and privacy is a continuously negotiated process in face-to-face interaction, online interactions make the case-by-case decision-making process difficult. Users rarely interact with each other synchronously. Instead, decisions of privacy, what to disclose and how, must be made a priori and explicitly" (2008, 111). In other words, as Nissenbaum's framework would indicate, although information continues to be tagged with very specific contexts, users have to determine a set of ex ante rules for determining how information should flow in and between contexts, and to do so before knowing any of the details about the information itself. In terms of contexts, privacy on Facebook is managed through settings, not ad hoc decisions.[12] Second, the abstraction involved in asynchronous, online social networking encourages a gap between a user's perceived audience and actual audience. As we will discuss in more detail below, users tend to significantly under-perceive the size and scope of the potential audience for their postings. College students, for example, tend to view their audience in terms of various peer groups, and not consider that administrators, faculty, family members and others might also be looking at the information.[13] Although offline gossip serves a loosely similar function, it is enormously magnified online (Solove 2007a). Third, although Facebook theoretically has a highly granular set of privacy settings, users do not appear to be taking advantage of them. On the one hand, the settings are byzantine, difficult to find, and hard to understand, thereby creating a significant barrier to users' ability to form or effectuate their privacy preferences. On the other hand, the structure of the system encourages a binarization of social relations into "friend" and "not friend," flattening out all of the nuances of face-to-face interactions.[14] Finally, as the following analysis will demonstrate, these general points need to be interpreted through the details of specific features of SNS program interfaces. Whether and how posting personal information online affects users who do so has a lot to do with the specific interfaces and policies of the sites where the information is posted. Offline, users can effectuate their privacy preferences by, for example, having a conversation behind closed doors. The architecture of physical spaces imposes considerable limitations on this process, however. Because the architecture of an online environment is a

[11] For an overview, see Strater and Lipford (2008); for another listing of risks, see Grimmelman (2009), 1164-78. One early study (Jones and Soltren 2005; cited in Solove 2007a, 197) found that the vast majority of Facebook users had not read the company's privacy policy, and almost two-thirds said they were not very concerned about privacy. Joinson (2008) records the first survey data in which a majority of respondents had modified their privacy settings. Lampe et al. (2008) also found a slim majority of users have changed their privacy settings, but staying with default settings was correlated with a decrease in the size and type of anticipated audience, suggesting that these users were not particularly savvy about privacy in general. Christofides et al. (2009) find evidence that the psychological processes that lead users to display information are different from those that lead them to control information—i.e., that there is no necessary tradeoff between the desire to control and display information. Livingstone (2008) finds that teenagers care deeply about privacy on SNS, but are frustrated in their efforts to manage the software's privacy controls.

[12] Of course, one can use rules for privacy settings offline, too ("never talk to a stranger"). Offline, however, rules can be more easily given exceptions ("it's ok to tell a stranger if you're hurt and need help"), and there are a variety of discrete but difficult-to-code ex ante decision strategies available other than rules. For a taxonomy, see Sunstein and Ullmann-Margalit (1999). As Bruno Latour puts it in a slightly different context, "no human is as relentlessly moral as a machine" (1992, 232).

[13] See Lampe et al. (2007), 168. Lampe et al. (2008) report that, although the number of college student users who thought that future employers might see their Facebook information was increasing, it was still quite low, and a large number of them viewed Facebook as a "student only" space. DiMicco and Millen (2007) found that a large number of (particularly younger) users of a corporate network on Facebook reported being unconcerned about coworkers seeing their profiles, since they viewed Facebook as outside work.

[14] For empirical evidence on these points, see Strater and Lipford (2008). For a similar suggestion about the loss of nuance on social networking sites, see Solove (2007a), 202. The difficulty in finding and effectuating privacy preferences undermines arguments to the effect that users have "chosen" to freely disseminate their information online. Livingstone (2008) reports that teenagers find this binarization on SNS frustrating.

function of the code that creates it, and because that code can be changed, the coding and interface of a site can make an enormous difference both in how much privacy users have, and how they experience their privacy.[15]

The precursors to social networking (blogging and webcams) demonstrated two phenomena—the risk of collateral privacy damage, and the plasticity of norms—and these phenomena are also evident in recent changes to Facebook. We will examine the case of collateral damage first, because its analysis is more straightforward than the plasticity of norms. In both cases, we will suggest that the difficulties lie in the tensions between the tendency of offline social contexts to be highly granular and specific, and the tendency of social network sites to mirror the more global context of the Internet, even as they substantially nuance that context through the emerging norms of online friendship. This granularity gap needs to be addressed through interface and design decisions that highlight, rather than obscure, the new ways that information travels online.

Collateral damage

In May 2007, Facebook introduced their Application Platform, allowing third party developers to create added functionality that links to a user's profile. These applications enhance the social experience on Facebook by allowing users to add additional content to their profiles, play games with their friends, share photos and other media, and much more. Applications have been extremely successful. Facebook reports that 70% of users interact with an application each month, with over five hundred thousand applications available. In order to complement a user's profile, Facebook allows applications to access most of the user's profile information, except for contact information. More disturbing, however, is that these applications are also allowed to access the same information for all of a user's friends. While this allows applications to incorporate information about a user's social spheres into their functionality, few need access to such a wide variety of information to do so (Felt and Evans 2008).

The privacy problems with such applications are easy enough to see. If I join a tennis club, I expect to tell them my name and address, as well as some information about my fitness level and maybe even my doctor's name or my birthday. I do not expect to share which books and movies I like, where I went to school and where I work, and what my religious and political affiliations are. And my friends have every reason to expect that I will not share the parallel information about them with the club. Doing so would be an egregious violation of well-entrenched norms of distribution. However, this is exactly what is happening with Facebook's third party applications. Worse still, this information sharing is largely invisible. Users are alerted with a simple message each time they install an application that both their own and their friends' information will be shared. However, this message is not very descriptive, and is easy to ignore as users are more focused on the task of using the application than on their privacy. There is an important parallel here to the attention paid to software license agreements that users must accept in order to use software. These are read by less than 2% of users (Good et al. 2007). As with such agreements, it seems that many users simply 'click through' these privacy notices, ignoring the one important piece of information that alerts them about giving away their information and the information of their friends to third parties.

Another serious violation to the norms of distribution is that the developers behind the applications are also largely invisible, obscuring the fact that the information is in fact leaving the confines of Facebook and not just going to the user's friends. A college student in Germany can create a quiz application that matches people with the beer they are most like, running it on his computer in his dorm. Users accessing this application will likely be completely unaware that their profile information, and their friends' profile information, can be accessed by a student in a dorm room in another country. While a third party application developer, such as the student in the dorm room in Germany, may be acting completely ethically and only accessing the profile information needed for the application, there is no guarantee of this, and little vetting process—anyone can become a Facebook application developer.

Both of these problems (the sharing of friends' information and the invisibility of the developer) are aggravated by the fact that users build their mental models of what applications can see and do based upon their regular and daily interactions, both with their friends and with Facebook. Applications run within the boundary of the site, giving users the impression that they are interacting with Facebook and others on Facebook. This effectively obscures the fact that they are also interacting with some third party server, such as an unsecured Windows machine in a student's dorm room in another country. The same user who installs a Facebook application would be likely to reject an email advertisement from an unknown person, encouraging her to install a similar application on her computer. The Facebook shell gives users a false sense of security. Additionally, the interactions with applications tend to be very personalized—a user takes a quiz through a Facebook application to be told what city she should live in or what cartoon character she most resembles. Users see that some of their information is used, often their name and

[15] For the point about software and architecture, see Lessig (2006) and MacKenzie (2006).

photo, and shared with friends who also use that application. And they can see their friends' names and photos in some applications. But since few applications actually make use of the majority of the profile information, users are not aware that all of the other personal information on their profile is potentially being accessed by those third party application developers. The result is that users have little understanding of the information they are sharing, with whom they are sharing it, and that they are responsible for sharing all of their friends' information as well (Besmer and Lipford 2010).

The invisibility of information flows presents a particular problem because when we do not know what is being done with our information, we have no ability to contest it. If the architecture and interface of Facebook essentially hides the amount of information that is shared to third parties, then there is little that a user can do about that sharing. Indeed, there is little that she can do to avoid the situation, other than decline to use the third party applications. The choice for users is binary: install the application and give full access to their own and their friends' personal information, or don't use the application at all.[16] We can imagine designing a number of different mechanisms that could provide greater transparency (and in some cases, greater control) of these information flows. The installation message could actually contain concrete examples of information that is shared: "Allowing NewCoolApplication access will allow it to pull most of the information from your profile, such as <your religion> and <your political affiliations>, your <##> photos, and all of the same information from everyone you have friended, including friends such as <friendA>, <friendB>, and <friendC>." The use of actual data and actual friend names could spur users to pause, read and reflect, rather than click through. Or, the user could be given the option to choose what information about herself and her friends she is willing to share.[17] Another method for increasing transparency and control could involve the use of warnings sent to friends anytime a person installs an application, allowing the friends to opt out of the information sharing.[18] Or, everytime an application asks for a piece of information from a user's profile, the user could be notified of that access.

Absent such transparency about sharing information, however, the relevant debates about privacy and sharing could never happen. This matters, because the acceptance of sharing such information with third parties would represent a significant change in norms of information distribution. Nissenbaum suggests that, in cases where it is necessary to decide between status quo norms and their change, there ought to be a strong presumption in favor of the status quo. Here, it is impossible either to make that presumption or to rebut it; a system that is engineered to hide deviations from distributional norms all but guarantees they will be violated. Ordinary people will make an ordinary assumption about norms: not about the content of those norms, but that they function in the new context of Facebook much like they function in the more familiar context of (say) joining tennis clubs. Having thus made this assumption, they will be unable to assess whether the difference in norms is one they want to endorse.[19] Similar sorts of issues emerge in the context of the News Feed, to which we now turn.

What's wrong with plastic norms?

Initially, if one wanted an update on one's Facebook friends, one had to go to their pages and read any updates they had posted. If one has very many friends, this is obviously an inefficient use of time. After all, at any given moment, only some of them will have posted updates, but all of them will have to be searched. An elegant solution to this problem borrows from blog and news aggregators: Facebook's "News Feed feature," abruptly introduced in 2006, gives users their friends' updates automatically, without the need to visit anyone else's page. To the

[16] There is an added sense of futility for the few users who really are aware of the policy that third party application developers get access to a user's friends' information—if my friend has already installed the application, the third party developer already has access to (at least some of) my data, so I have to weigh the benefit of using a service from a developer who already has my data against the privacy concerns of giving away both more of my own data and the data of my friends, whose information might not already be accessible to that developer.

[17] Facebook has recently proposed such a modification, where users will expressly indicate consent for each category of information that a particular application would like to access. This change is in response to an investigation by the Office of the Privacy Commissioner of Canada (2009). While this is a step in the right direction, we suspect that users will still not be fully aware of the implications of consent—namely that information is flowing off of the site to a third party and that friends' information is also shared.

[18] While there are currently privacy settings to limit the amount of information given to a friend's applications, with no understanding of their purpose, there is little incentive to use them, see Besmer and Lipford (2010).

[19] For a listing of cognitive mistakes about privacy risk that Facebook's imitation of offline friendship encourages in its users, see Grimmelman (2009), 1160-64. Although we will not pursue the point here, we think that it would be possible to develop a more thorough normative grounding of the above argument with reference to Kant's formula for humanity. In Kant's central example of deceit, the point is that the other person is literally unable to assent to the deceit, because the deceit hides the actual bargain being offered. Similarly,companies that hide the degree to which they share information about uninvolved parties might be vulnerable to the charge that they render "consent" to privacy practices meaningless.

surprise of Facebook management, there was a strong backlash against the News Feed's rollout. The questions raised by the News Feed are different from Applications, because the flows of information with News Feed are by and large visible. The openness of norm creation and change that Applications make impossible is readily available with the News Feed. We will first discuss issues related to the introduction of the feature, and then to the feature itself.

The backlash was in part generated by a sense that the News Feed violated the privacy of users. In the terms of Nissenbaum's framework, the analogy would be to the case of public records being placed online.[20] For public records, she suggests, the privacy problem is one of the scope of the distribution of the information. It is not that there is any more information publicly available; it is that it is a lot easier to get to the information. It is no longer necessary to travel to a (possibly remote) local courthouse. When an activity is easier, more people will do it. Thus, for example, idle curiosity about 'who owns that house, and what they paid for it' becomes a topic that one might actually pursue online, when the same topic would be too much like work to pursue offline. As a result, people have less *de facto* privacy. Facebook similarly increased the efficiency of accessing information that was already in principle accessible; in so doing, it reduced the *de facto* privacy of its users. As danah boyd puts it, "Facebook made what was previously obscure difficult to miss (and even harder to forget). Those data were all there before but were not efficiently accessible" (2008, 15). She adds:

> Information is not private because no one knows it; it is private because the knowing is limited and controlled. In most scenarios, the limitations are often more social than structural…. There is an immense gray area between secrets and information intended to be broadcast as publicly as possible. By and large, people treated Facebook as being in that gray zone …. When snippets and actions were broadcast to the News Feed, they were taken out of context and made far more visible than seemed reasonable. In other words, with News Feeds, Facebook obliterated the gray zone (2008, 18).[21]

The switch to News Feeds did violence to users' norms of distribution, and it seems quite clear that the company could have, and should have, managed the introduction of the feature much better than it did. As Nissenbaum puts the general point, we have good reason to favor the status quo in privacy situations on the grounds that they reflect established patterns of social interaction; "common practices are understood to reflect norms of appropriateness and flow, and breaches of these norms are held to be violations of privacy" (2004, 127).

Not long after that, however, a funny thing happened: the backlash dissipated, and users came to accept the News Feed, even beginning to manipulate postings to their advantage.[22] How are we to understand this apparent lability of users' privacy norms? As a general point, one should note that the novelty of social networking sites means that norms for behavior on those sites will be relatively unsettled, without the stabilizing weight of long tradition. In addition, as noted above, one difference between being online and being offline is that, online, the basic architectural features of one's context are both more fluid and more easily changed. As Nissenbaum points out, privacy norms are going to depend on, and change with, context, and "these context-relative qualifications can be built right into the informational norms of any given context" (2004, 138). Together, these imply that we should expect to see a greater range of privacy norms online, and that those norms could change in ways not easily paralleled offline. In this specific case, the ability of users to embrace News Feed seems to stem from three factors. First, the switch to News Feed essentially changed the distributional context of status updates. Once users adjusted to the different context, they were able to adjust their behavior accordingly. Second, having adjusted to the new context, they saw benefits to using it. Finally, the privacy preferences in question were not particularly intense, because News Feed did not implicate any perceived core privacy issues: users were already sharing information widely; the question was the scope of its distribution.

Since users have embraced News Feeds, and since Facebook's privacy norms appear to be somewhat fluid, it is appropriate to revisit the privacy complaint while taking the new norms into account. As Nissenbaum points out, and as we have illustrated in the case of Applications, new

[20] There are limits to the analogy. Facebook is an opt-in system (with an opt-out mechanism) in a way that public records like housing ownership information are not. If I buy a house, that information is automatically public, and there is nothing I can do about it. For somebody to get my Facebook updates, I have to have shown enough interest in them at some point to friend them. Friending might be casual and so reflect only a *de minimis* level of interest, but it is also clearly not the same thing as automatically public. So too, if I want individuals not to get my Facebook information, I can un-friend them; there is no analogous procedure for legal records.

[21] For compatible analysis, see Solove (2007a), 170 and 198.

[22] For a critical discussion, see boyd (2008); for evidence of acceptance, see Joinson (2008, 1031). One topic which we will not pursue here, but which bears at least a mention, is the extent to which the switch to the News Feed encourages users to spend more time on the site posting updates. Since the value of the site is generated by the amount of time users spend on it, the switch to the News Feed could be read as a (successful) effort to get users to contribute more free labor to Facebook's coffers. This worry is pursued in the case of the Internet more generally in Terranova (2004), 73–97.

privacy practices can challenge our status-quo presumption, and force us to reflect on status quo norms "in terms of how well they promote not only values or goods internal to a given context, but also fundamental social, political, and moral values" (2004, 128). It is therefore appropriate to ask: is there a per se privacy problem with the News Feed, and how could Facebook best mitigate any potential privacy problems? As long as one stays with the news or blog aggregator analogy, the News Feed is clearly an enormous improvement in efficiency, since it saves having to individually visit potentially dozens of pages. Not only that, there is no obvious, inherent privacy problem with blog postings or news aggregation (as the Washingtonienne example shows, of course, that does not mean there are no problems possible). However, privacy and efficiency are often in tension with one another, and we should ask what the inefficiency of the old system enabled. Like news stories, updates are generally intended for public consumption. Of course, that only friends get updates significantly narrows the audience, but the update is still not like a personal email. However, there does seem to be a difference between getting news automatically and having to seek it out; this difference is what fueled the initial backlash against newsfeeds. In particular, the updates go to all of one's friends—whether or not they are interested enough to seek out the information themselves. The updates might thus seem like performance-enhanced versions of family holiday letters. Every day—sometimes many times a day—a Facebook user learns that friend Joe has been succeeding at school and that cousin Mary is really enjoying her soccer team.

A more significant difference occurs on the other end, as newsfeeds do not just automate the process of receiving updates; they automate the process of sending them. Here a significant difference from holiday letters emerges: the process of sending those letters is inefficient, in that it involves individually addressing and stuffing envelopes for every update. If the letter says something that a recipient might find offensive, he gets a card but not the letter. If the letter is going to someone that a person no longer cares to keep up with, she might not send it, and might even remove the person from her list. In other words, being offline in this case introduces a necessary pause between composing the letter and sending it, and this encourages the list itself being subject to at least a casual vetting procedure. No such procedural encouragement exists for Facebook updates. Nothing in the process of sending updates encourages one to even notice who the full audience is, much less to assess who should be in it. Updating friends lists is a separate procedure; Facebook doesn't hide the option, but the design of the system doesn't highlight it either.

Other aspects of the Facebook experience converge to facilitate indiscriminate disclosures. First, Facebook users tend to operate with a "shrinking perceived audience." That is, they initially begin by friending a large number of people, and assuming that everything they say is more or less public. Over time, and as their active circle of friends narrows, they tend to forget about the earlier friends and strangers who can potentially access their information because that access seems unlikely or infrequent. Their postings then reflect their perception that those they interact with are constitutive of their audience, rather than all those who could potentially view their information.[23] In an opt-in world of updates, this would be far less significant than in an opt-out one, since the act of looking for an update would be evidence that the person was still in one's circle (if not the perceived one). Opt-out updates, on the contrary, encourage the gap between the perceived and actual circle of friends. Second, users' shrinking perceived audience has a temporal feature in that most users tend not to think about how their updates remain available more or less permanently. The News Feed format also encourages this thought, since updates show up on one's list, only to be displaced shortly thereafter by other updates. One's experience, then, is of ephemeral news postings, not a permanent record. But the record is nonetheless permanent by default; it is possible to go back and view all of a person's updates over a period of several years unless she deletes them. Third, Facebook interactions *feel* bilateral, even when they're not. For example, 'wall-to-wall' communications allow one to exchange messages with a single friend asynchronously in what feels and looks like a private space, but which is in fact visible to others. Finally, the privacy settings further discourage user self-awareness. Facebook recognized the difference between opt-in and opt-out updates, and added privacy settings allowing a user to control what types of information will be broadcast, and to whom. Unfortunately, these settings are no easier to use than Facebook's other privacy features.

These problems have as much to do with Facebook's interface as with the underlying architecture of the News Feed. That matters because the normative status of "friendship" on the site is constantly evolving, as the technology of social networking enables entire types of interactions that are not available offline.[24] In this respect,

[23] For further discussion and data supporting this point, see Strater and Lipford (2008).

[24] For a discussion of some of the complexities involved in online "friendships," see boyd (2006). Through her discussion of Myspaces's "Top 8" feature, boyd usefully emphasizes the extent to which the meaning of friendship in SNS is partly a function of the design of the sites; as she puts it, "friending supports pre-existing social norms, yet because the architecture of social network sites is fundamentally different than the architecture of unmediated social spaces, these sites introduce an environment that is quite unlike that with which we are accustomed. Persistence, searchability, replicability, and invisible audiences are all properties that participants must negotiate when on

social networking picks up where blogging, webcamming, and other kinds of online interaction leave off. Since, to a much greater extent than is the case with blogs and webcamming, social networking is mediated through the highly structured and externally provided environment of sites like Facebook, changes in those sites' interfaces can significantly affect how people behave online, and those behavioral changes feed back into the norms that guide them. For example, most of us do not have access to mass media outlets with which to simultaneously update a large body of mostly casual acquaintances on our every move, and so the move to the News Feed accentuates at least one way in which Facebook social relations are significantly different from their offline counterparts. At the same time, the interface facilitates the tendency of users not to notice this difference over time. For this reason, as with Applications, the change of norms is interesting not just as a privacy issue, but, inseparably from that, one of design. In the case of the News Feed, the issues emerge even in implementing a feature that does not seem obviously objectionable, and where the flows of information enabled by that feature are not thoroughly hidden.

There are a variety of ways that the interface could be modified to deal with the situation. All of them turn on the thought that, when changes in context generate changes in flows of information, maintaining privacy requires drawing people's attention to the changes in question. The software could, for example, make updates unavailable after a certain amount of time. The interface would then be structured on the basis of users' expectations of it. The software could also prominently provide a "delete old updates" button, which would remind users that their updates are permanently on record, unless they take explicit counter-measures. The software could also default to deleting updates, requiring users to "save old updates." Alternatively, the interface could be changed from the point of view of the reader: attached to each update could be a "view all of Mary's updates" option, which would subtly remind users that the same option applies to their own updates. It could even be designed to send a notice to users: "Mary has just looked at all of your updates." This is not to endorse any one of these design options; the point is to underscore that each of these design features embeds normative preferences about the distribution of information and how they develop. Automatic deletion of old updates, for example, would move Facebook's community closer to an offline small town, where gossip travels quickly but is imperfectly remembered. Retention of updates, combined with reminders to that effect, would further encourage users to view their identities online as constructed and performative, moving Facebook closer to the world of Jennicam.[25] This encouragement would be magnified even more by a "Mary has just looked at all of your updates" option; users would increasingly view their Facebook identities as subject to constant surveillance, and modify them accordingly. If I knew that Mary always looked at all my updates, I might update with her in mind. More generally, users would be encouraged to have updates resemble the news wires from which the model was taken. Whether or not this is a "good" model for online friendship on a social networking site could then be discussed by users who were encouraged to form a mental model that accurately reflected the various flows of information on the site.

Conclusion

It is by now generally accepted by all but the most libertarian that technologies and their interfaces can facilitate some values and behaviors at the expense of others. For example, although a door that closes too slowly wastes heat, a door that closes too quickly "discriminates" against the disabled and delivery personnel. An intellectual property regime that grants strong proprietary rights and few exceptions favors mass media over independent producers of speech. Even the placement of advertising on portal sites matters. Recent work comparing Facebook with other SNS

Footnote 24 continued
social network sites." One might object that online and offline interactions are so different that "friendship" cannot apply to both. This seems to us to assume what is to be proven, namely, what "friendship" means. In lieu of a detailed discussion, we will make two brief points: (a) the term is in common use in both contexts; this use needs to be understood before discarding it; and (b) that Facebook is largely used to support and develop offline social networks suggests that insisting on a sharp, conceptual rupture between on and offline relations may only serve to obscure similarities in the name of highlighting differences. In any case, we are less concerned with the terminology than the social relations.

[25] This is not to claim that identity on Facebook is not tethered to offline identity; it clearly is. However, insofar as Facebook supports offline interactions by allowing users to maintain and deepen offline relationships, the site allows users significant latitude in which aspects of their identities they emphasize. This happens both because many of the usual social cues through which we reveal ourselves (even unconsciously) do not port well to an online environment (see Cocking 2008 for related thoughts on privacy in this regard), and because SNS allow the asynchronous presentation of substantial amounts of information that would be unavailable to a casual, offline relationship. Those who interact with SNS users both off and online will then have to translate between how individuals present themselves in different contexts. This process of translation, which of course happens any time that two individuals encounter each other in different social contexts, sets a limit to identity construction online, insofar as it has to be possible to credibly translate between the multiple self presentations. For evidence that Facebook users do manage their identities online, see DiMicco and Millen (2007); for identity construction by teenagers on Myspace, see boyd (2007); for reflection on some of the ethical implications of all this, see Fleckenstein (2008).

suggests that the architectural features of those sites will tend to facilitate certain kinds of social interactions and not others.[26] The preceding analysis underscores the important relations between privacy norms and application and interface design. Two points bear emphasis. First, at the core of the privacy issues discussed here is a gap between the granularity of offline social contexts and those on Facebook. One of the greatest strengths of Nissenbaum's contextual integrity framework is that it highlights the importance of social contexts to privacy norms. The application of Nissenbaum's work to Facebook underscores the extent to which the granularity differences are driving many of the surface-level privacy issues. Facebook, as we have argued, does not adequately reflect this insight. As a result, offline contexts are more easily differentiated and kept separate than those on Facebook. To the extent that Facebook users use the program to develop offline social ties, this granularity gap will inevitably produce conflicts.

Second, this gap shows that part of what is at stake is the definition of what it means to be a "friend" on Facebook. The switch to the News Feed shows that the norms for online friendship are in a state of considerable flux. While no-one would claim that a Facebook friend is the same as an offline friend, the fact that there are many overlapping cases means that the shared vocabulary sometimes leads us to a misperception of sameness. In fact, the concepts of a 'friend' on Facebook and an offline friend may be diverging even more, making this difference bigger. The News Feed makes updates a lot like blogging, and Applications put users in the position of treating their friends' information as commodities to exchange for access to whatever the application in question does. If Facebook norms for friendship move further away from those offline, we would expect that users' would model their online behavior accordingly. While few of them might endorse the sense of complete publicity that Jennicam embraced, publicity at nearly that level is enabled by the combination of the News Feed and a user who posts constant updates.

These points suggest that many of the privacy issues on Facebook are principally design issues, by which we mean that they can be addressed through better program design.[27] As Nissenbaum points out, users base their behavior on the expected norms of the perceived context, and that perception on Facebook is influenced by the design of the site. In particular, Facebook needs to do a better job of making the flows of information on the site transparent to users. As noted above, Facebook users face a tradeoff between disclosure and control of information, and there is reason to think both that they do not perceive the two as trading off with one another, and that, to the extent that they do perceive the tradeoff, they misunderstand how risks are allocated. Good program design can push users in the direction of understanding these tradeoffs, so that they can make usage decisions intelligently. On the one hand, users need to be aware of what, exactly, might happen with their and their friends' information, in order to make informed decisions about how to share that information. These decisions will then form a part of what it means to be a Facebook friend. On the other hand, insofar as the cognitive model users bring to Facebook is offline friendship, they will need extra reminders that the ways they handle information and privacy offline do not port directly to Facebook. Some of these design decisions might be remedied at the architectural level, in the sense that an application might be restricted to only the information it needs to operate. Other design improvements would be at the interface level, in the form of helping users form correct cognitive models of information flows online. For example, we suspect that many users of Facebook would not endorse any notion of "friend" that included dispensing all of one's personal information to unknown third parties. Users need to be aware of the outward flow of this information so that their consent, if they give it, can be informed. Similarly, users need to be reminded that their updates are going to be seen by people whom they have forgotten are part of their audience. Whatever else Jennicam did right or wrong, the site had the merit of making it impossible to ignore the flow of information out of Ringley's apartment and onto the Web, and of effectively posing that flow of information as a moral question. Since the meaning of their friendship is at stake, Facebook friends deserve no less transparency about information flows, and no less opportunity to develop their own privacy norms.

[26] For the doors, see the classic discussion in Latour (1992); for intellectual property regimes, see Benkler (2003); for portals, see Introna and Nissenbaum (2000); for the extension to SNS, see Papacharissi (2009).

[27] While we have offered a number of different examples of designs that could increase transparency (and control) of information flows, our goal is not to prevent all information sharing (which would be counter to the whole point of social networking), but rather to promote interface designs which give users the freedom to make informed choices about what they share and with whom. We have further discussed design guidelines and example solutions in Lipford et al. (2009). Dwyer and Hiltz (2008) also discuss the need to improve

Footnote 27 continued
the design of privacy mechanisms in online communities such as Facebook. Grimmelman (2010) proposes an analogy to product safety law, where vendors are responsible for designing products to promote consumer safety by, for example, preventing dangerous misuse, making consequences of actions predictable, working with and not against consumer expectations, and by making sure that product warnings are of sufficient relevance and quality. He finds Facebook lacking on many of these counts.

References

Alexa. (2010). Facebook.com—site info from Alexa. March, at: http://www.alexa.com/siteinfo/facebook.com.

Bailey, J. (2009). Life in the fishbowl: Feminist interrogations of webcamming. In *Lessons from the identity trail: Anonymity, privacy and identity in a networked society*. Oxford: OUP, pp. 283–301.

Benkler, Y. (2003). Through the looking glass: Alice and the constitutional foundations of the public domain. *Law and Contemporary Problems, 66*, 173–224.

Benkler, Y. (2006). *The wealth of networks: How social production transforms markets and freedom*. New Haven: Yale UP.

Besmer, A., & Lipford, H. R. (2010). Users' (Mis)conceptions of social applications. To appear, *Proceedings of Graphics Interface*.

Binder, J., Howes, A., & Sutcliffe, A. (2009). The problem of conflicting social spheres: Effects of network structure on experienced tension in social network sites. *Proceedings of CHI 2009*, Boston: ACM Press, pp. 965–974.

boyd, d. (2006). Friends, Friendsters, and MySpace Top 8: Writing Community Into Being on Social Network Sites. *First Monday* 11:12, December. http://www.firstmonday.org/issues/issue11_12/boyd/index.html.

boyd, d. (2007). Why youth (heart) social network sites: The role of networked publics in teenage social life. In D. Buckingham (Ed.) *MacArthur foundation series on digital learning—youth, identity, and digital media volume*. Cambridge, MA: MIT Press, pp. 119–142.

boyd, d. (2008). Facebook's privacy trainwreck: Exposure, invasion and social convergence. *Convergence, 14*(1), 13–20.

Brey, P. (2000). Disclosive computer ethics. *Computers and Society* (Dec.), pp. 10–16.

Christofides, E., Muise, A., & Desmarais, S. (2009). Information disclosure and control and facebook: Are they two sides of the same coin or two different processes? *Cyberpsychology and Behavior, 12*(30), 341–345.

Clark, A. (2003). *Natural-Born cyborgs: Minds, technologies, and the future of human intelligence*. Oxford: OUP.

Cocking, D. (2008). Plural selves and relational identity. In J. van den Hoven & J. Weckert (Eds.), *Information technology and moral philosophy* (pp. 123–141). Cambridge: CUP.

Cutler, J. (2004). Senator sacked me over tales of congress, *The Guardian (UK)*, June 2, 2004, at http://www.guardian.co.uk/world/2004/jun/02/usa. Accessed 5/09.

DiMicco, J., et al. (2008). Motivations for social networking at work. In *Proceedings of CSCW'08*. San Diego: ACM Press, pp 711–720.

DiMicco, J. M. & Millen, D. R. (2007). Identity management: Multiple presentations of self in facebook identity. In *Proceedings of GROUP'07*. Florida: ACM Press, pp. 383–386.

Dwyer, C., & Hiltz, S. R. (2008). Designing privacy into online communities. In *Proceedings of Internet Research 9.0*, October 2008.

Ellison, N., Steinfeld, C., & Lampe, C. (2007). The benefits of Facebook 'Friends:' Social capital and college students' use of online social network sites. *Journal of Computer Mediated Communication,* 12(4), article 1, http://jcmc.indiana.edu/vol12/issue4/ellison.html.

Felt, A., & Evans, D. (2008). Privacy protection for social networking APIs. In *Proceedings of Web 2.0 Security and Privacy 2008*.

Fleckenstein, K. S. (2008). Cybernetics, ethos, and ethics. In L. Worsham & G. A. Olson (Eds.), *Plugged in: Technology, rhetoric and culture in a posthuman age* (pp. 3–23). Cresskill, NJ: Hampton Press.

Gelman, L. (2009). Privacy, free speech, and 'blurry-edged' social networks. *Boston College Law Review, 50*, 1315–1344.

Good, N. S., Grossklags, J., Mulligan, D. K., & Konstan, J. A. (2007). Noticing notice: A large-scale experiment on the timing of software license agreements. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. San Jose: ACM Press, pp. 607–616.

Grimmelman, J. (2010). Privacy as Product Safety. *Widener Law Journal* (forthcoming), available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1560243.

Grimmelman, J. (2009). Saving Facebook. *Iowa Law Review, 94*, 1137–1206.

Gross, R., & Acquisiti, A. (2005). Information revelation and privacy in online social networks. In *Proceedings of WPES'05*. Alexandria, VA: ACM Press, pp. 71–80.

Hashemi, Y. (2009) Facebook's privacy policy and its third-party partnerships: Lucrativity and liability. *Boston University Journal of Science and Technology Law,* 15, 140–161.

Hull, G. (2009). Overblocking autonomy: The case of mandatory library filtering software. *Continental Philosophy Review, 42*, 81–100.

Introna, L. D., & Nissenbaum, H. (2000). Shaping the web: Why the politics of search engines matters. *The Information Society, 16*(3), 169–185.

Jimroglou, K. M. (1999). A camera with a view: JenniCAM, visual representation, and cyborg subjectivity. *Information, Communication & Society, 2*(4), 439–453.

Joinson, A. N. (2008). 'Looking at,' 'Looking up,' or 'Keeping up with' People? Motives and uses of Facebook. In *CHI 2008 Proceedings: Online Social Networks*. pp. 1027–1036.

Jones, H., & Soltren, J. H. (2005). Facebook: Threats to privacy, available at: http://groups.csail.mit.edu/mac/classes/6.805/student-papers/fall05-papers/facebook.pdf. Accessed 5/09.

Lampe, C., Ellison, N. B., & Steinfield, C. (2006). A Face(book) in the crowd: Social searching vs. social browsing. *Proceedings of CSCW'06*. Alberta: ACM Press, pp. 167–170.

Lampe, C., Ellison, N. B., & Steinfield, C. (2007). A familiar Face(book): Profile elements as signals in an online social network. In *Proceedings of CHI'07*. San Jose: ACM Press, pp. 435–444.

Lampe, C., Ellison, N. B., & Steinfield, C. (2008). Changes in use and perception of Facebook. In *CSCW'08*. San Diego: ACM Press, pp. 721–730.

Latour, B. (1992). Where are the missing masses? The sociology of a few mundane objects. In W. E. Bijker & J. Law (Eds.), *Shaping technology/building society* (pp. 225–258). Cambridge, MA: MIT Press.

Lessig, L. (2006). *Code and other laws of cyberspace, Version 2.0*. New York: Basic Books.

Lipford, H. R., Hull, G., Latulipe, C., Besmer, A., & Watson, J. (2009). Visual flows: Contextual integrity and the design of privacy mechanisms on social network sites. In *Proceedings of the Workshop on Security and Privacy in Online Social Networking, IEEE International Conference on Social Computing* (SocialCom), August 2009.

Livingstone, S. (2008). Taking risky opportunities in youthful content creation: Teenagers' use of social networking sites for intimacy, privacy, and self-expression. *New Media Society, 10*, 393–411.

MacKenzie, A. (2006). *Cutting code: Software and sociality*. New York: Peter Lang.

Matthews, S. (2008). Identity and Information Technology. In J. van den Hoven & J. Weckert (Eds.) *Information technology and moral philosophy*. Cambridge: CUP, pp. 142–160.

Myskja, B. K. (2008). The categorical imperative and the ethics of trust. *Ethics and Information Technology, 10*, 213–220.

Nielsen, W. (2009). Time spent on Facebook up 700%, but MySpace Still Tops for Video, at http://blog.nielsen.com/nielsenwire/online_mobile/time-spent-on-facebook-up-700-but-myspace-still-tops-for-video/. Accessed 6/09.

Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review, 79*, 101–139.

Office of the Privacy Commissioner of Canada. (2009). Facebook agrees to address Privacy Commissioner's concerns, August 27, 2009. at http://www.priv.gc.ca/media/nr-c/2009/nr-c_090827_e.cfm. Accessed 11/09.

Papacharissi, Z. (2009). The virtual geographies of social networks: A comparative analysis of Facebook, LinkedIn and ASmallWorld. *New Media Society, 11*, 199–220.

Solove, D. J. (2007a). *The future of reputation: Gossip, rumor and privacy on the internet*. New Haven, CT: Yale UP.

Solove, D. J. (2007b). 'I've Got Nothing to Hide' and other misunderstandings of privacy. *San Diego Law Review, 44*, 745–772.

Strater, K., & Lipford, H. R. (2008). Strategies and struggles with privacy in an online social networking community. In *Proceedings of the 22nd British HCI Group 2008*. Liverpool UK: ACM Press, pp. 111–119.

Sunstein, C. R., & Ullmann-Margalit, E. (1999). Second-order decisions. *Ethics, 110*, 5–31.

Technorati. (2008). *State of the blogosphere*, at http://technorati.com/blogging/state-of-the-blogosphere/. Accessed 5/09.

Terranova, T. (2004). *Newtork culture: Politics for the information age*. London: Pluto Press.

Wesch, M. YouTube and you: Experiences of self-awareness in the context collapse of the recording webcam. ms. on file with authors (Forthcoming).