HUMAN–COMPUTER INTERACTION, 2006, Volume 21, pp. 319–342
Copyright © 2006, Lawrence Erlbaum Associates, Inc.

Collective Information Practice: Exploring Privacy and Security as Social and Cultural Phenomena

Paul Dourish
University of California, Irvine

Ken Anderson
Intel Corporation

ABSTRACT

As everyday life is increasingly conducted online, and as the electronic world continues to move out into the physical, the privacy of information and action and the security of information systems are increasingly a focus of concern both for the research community and the public at large. Accordingly, privacy and security are active topics of investigation from a wide range of perspectives—institutional, legislative, technical, interactional, and more. In this article, we wish to contribute toward a broad understanding of privacy and security not simply as technical phenomena but as embedded in social and cultural contexts. Privacy and security are difficult concepts to manage from a technical perspective precisely because they are caught up in larger collective rhetorics and practices of

Paul Dourish is a computer scientist whose work lies primarily at the intersection of computer science and social science, with research interests in Computer-Supported Cooperative Work and Ubiquitous Computing; he is Professor of Informatics and Associate Director of the California Institute for Telecommunications and Information Technology at the University of California, Irvine. **Ken Anderson** is an anthropologist with research interests in identity, culture, and technology; he is Senior Researcher and Manager in Intel Research at Intel Corporation.

CONTENTS

risk, danger, secrecy, trust, morality, identity, and more. Reductive attempts to deal with these issues separately produce incoherent or brittle results. We argue for a move away from narrow views of privacy and security and toward a holistic view of situated and collective information practice.

1. INTRODUCTION

It is scarcely a novel observation that information technology is colonizing ever-larger regions of our lives. Laptops, PDAs, MP3 players, cell phones, and even flash memory devices attached to keyrings are increasingly not just ubiquitous but essential for the conduct of everyday life. We are faced with a proliferation of computational devices that are ever more powerful computationally and accessible economically, and devices from home appliances to hotel doorknobs have started to exhibit digital capabilities.

As storage capacities rise and as more powerful computational engines make it possible to search, process, and filter huge volumes of information in ways not previously envisioned, we develop new ways of relating to information. We begin to question, for example, whether it is worth the effort of ever deleting information; easier, perhaps, to store everything on the off-chance that one day it will be needed. Knowing that any digital information we generate will likely be stored and indexed transforms how we communicate and present ourselves (Grudin, 2001). Analog information (e.g., security camera video) is increasingly available digitally and amenable to digital information processing, transforming the nature of surveillance and the information environment. Newer monitoring technologies such as RFID tags—once used largely to track farm animals—are cheap enough that retailers such as Walmart can insist that their suppliers include them in all products, easing inven-

tory management but also enabling new forms of electronic monitoring and profiling (Want, 2004).

There is a sense running through these developments that our new digital capabilities outstrip our understanding of how to use them. One of the significant debates about rapid technology development concerns the problem of privacy, which has become a major focus for both academic and popular attention (e.g., Agre & Rotenberg, 1997; Brin, 1999; Schneier, 2000; Westin, 1968). Ubiquitous computing, with its emphasis on capturing aspects of personal and collective "context" or operating in a pervasive sensing environment, is perhaps one area in which these issues are particularly visible, and in this area, privacy is widely acknowledged as a major research issue (e.g., Boyle & Greenberg, 2005; Langheinrich, 2002; Palen & Dourish, 2003). Accordingly, a number of system designers and developers have begun to focus on privacy as a key element of information system design, recognizing that, like usability, privacy concerns are not something that can be retrofitted to technologies but are fundamental to their structure and usage models. Despite these developments, however, there is little sense that privacy problems are being resolved. They seem stubbornly persistent.

The argument that we wish to develop here is that one reason for these problems is that, although authors such as Agre (1994), Clement (1994), and Clark (1988) have attempted to place technological privacy and security concerns within social and cultural contexts, technological developments have been slow to incorporate the understandings that they provide. Our goal then is to contribute to this broader exploration of privacy and security and to show how this can be connected to technological and design discussions. In particular, we propose that design attention should focus on privacy and security as aspects of "collective information practice" rather than as individual technical states.

In approaching this problem, we are explicitly attempting to bridge intellectual traditions. As a computer scientist with a deep concern for the social systems within which information technology is embedded and a cultural anthropologist with a fundamental orientation toward technology design, we bring to this investigation the belief that social and technical are irremediably recursively linked and that any solution to the privacy conundrum must start from this basis. It should be noted, too, that we come at the questions of privacy from a Western cultural context, and indeed many of the examples that we draw on share these cultural roots. Even in localized contexts, cultural practices around privacy are remarkably varied (Altman, 1977). This said, our attention is focused on the consequences for analysis and design in human–computer interaction (HCI), rather than more broadly.

Because both *privacy* and *security* are terms of art within the domains of information technology and HCI design, it is worth making clear what we

mean by them in this article. In the information technology domain, security describes, first, the ways in which information systems may be vulnerable to a range of conscious attacks in which the effectiveness of the system or the integrity of its operation might be compromised and, second, the technical countermeasures that can be used to detect, respond to, and protect a system from these attacks. Privacy, on the other hand, is something of a catch-all term that refers primarily to the ways in which individuals or organizations might lose control of access to information and to individuals. It embodies a number of different aspects, including solitude (the right to be left alone), confidentiality (the right to control one's own information), and autonomy (the right to control self-presentation) (Gavison, 1980).

Privacy, then, is generally approached as a social consideration, whereas security is seen as a technical concern. The relation between them is that security technologies might provide mechanisms by which privacy can be ensured. Here, however, we wish to take a slightly different perspective, which depends on a social reading of security. We read security here as the state of being free from danger; technological "security mechanisms" are deployed as means to ensure this state. Risks to privacy (solitude, confidentiality, autonomy), then, are among the various risks against which we might wish to be secure. Privacy and security are not synonymous but are strongly related and, as we outline, have often been approached from the same analytic perspective.

Our goal in this article is not to offer a new framework for design but to understand the social and cultural practices that lie beneath technical specifications of privacy and security and to explore their implications for how we understand privacy and security in everyday technological settings. In particular, we wish to argue four points:

First, that the dominant model of privacy and security–as "economic rationality" dominated by the logic of the cost–benefit trade-off (later)—neglects certain critical aspects of these concerns as enacted social practices. Instead, we suggest a focus on the practical and discursive elements of privacy and security, which asks not what privacy and security *are* but rather what privacy and security *do*.

Second, that formulations of privacy and security must, implicitly or explicitly, draw on or respond to models of risk and danger and that these have their origin not simply in the physical world but in the social and cultural encounters with that world. Formulations of risk and danger are ways in which distinctions between acceptable and unacceptable behavior are labeled.

Third, that flows of information also serve as markers of social boundaries, providing a means to negotiate, demonstrate, and sustain patterns of identity, membership, and affiliation in social groups. Secrecy and trust demonstrate the strategic use of information in this way and turn our attention away from information and toward the uses to which information is put.

Fourth, and consequently, that we may be able to engage in more effective design interventions by moving our focus away from information and its regulation and looking instead at what we term collective information practice—for ways in which social action is sustained and reproduced through the formulation and flow of information.

2. HCI PERSPECTIVES ON SECURITY AND PRIVACY

Our primary concern is with the ways in which privacy and security are understood as aspects of interaction between people and systems and, therefore, as topics for HCI design attention. We begin, then, by briefly exploring the ways in which these topics have been addressed within the HCI community.

The spread of the Internet as a primary platform for end-user computing has brought with it a number of concerns over system security (e.g., defense against computer viruses) and personal privacy (e.g., the safety of credit card information in e-commerce transactions, or commercial tracking of personal browsing behavior through "cookies"). These topics have been matters of concern for researchers in HCI (e.g., Ackerman, Cranor, & Reagle, 1999; Adams & Sasse, 1999; Millett, Friedman, & Felten, 2001). Although sophisticated security technologies are available, everyday practice seems continually to fall short of the levels we can achieve, not least because researchers in information security have generally aimed their solutions at system administrators rather than end-users, resulting in usability problems for everyday computer use (e.g., Good & Krekelberger, 2003). Recently, in response to this situation, a distinct community has begun to form whose central concern is the usability of security (e.g., Patrick, Long, & Flinn, 2003). Although it would take more space than we have here to review this work, we try to provide an overview of the major trends.

Broadly, we can distinguish three programs of work within this domain: empirical examinations of security practice, empirical studies of the usability of current and potential security technologies, and the design of novel mechanisms for managing privacy and security in end-user settings.

The first program concerns empirical investigations of the ways in which security and privacy are managed in real-world settings. For example, Angela Sasse and colleagues at University College London have been engaged in an extensive exploration of the interactions between security and usability, reported in a wide range of publications. Some of these have been predominantly explorations of security practice in real-world settings (e.g., Adams & Sasse, 1999; Weirich & Sasse, 2001), some have reported on empirical investigations of the effectiveness of particular techniques (e.g., Brostoff & Sasse, 2000; Riegelsberger et al., 2003), and others have developed new design models for secure interactive systems (Brostoff & Sasse, 2001; Flechais et al., 2003). An im-

portant element of this work, which concerns us here, is the social embedding of privacy practice; the design models that have emerged from these empirical investigations have repeatedly demonstrated the influence of social and organizational setting on security attitudes and behaviors. This work exemplifies a broad program of empirical studies examining privacy and security practices in real-world settings, including studies of instant messaging (Patil & Lai, 2005) and location-based services (Consolvo et al., 2005).

Whitten and Tygar's (1999) study was the first (and probably most influential) in a series of explorations of the usability of security software systems, a second program of work. Through an experimental evaluation of the use of public key encryption and digital signatures in an e-mail tool, they convincingly demonstrated that usability is a security issue and that problems understanding how to use security features result in people engaging in insecure information behaviors. The fact that security solutions are inherently complex and that security is a pervasive aspect of system design only serve to exacerbate the problem. This last point is particularly ironic; just as usability specialists have long argued that usability cannot be "grafted on" to a design once it is complete, security specialists have made just the same observation about security. This work has inspired a range of studies of the usability of security mechanisms (e.g., Brodie, Karat, & Karat, 2005; Millett et al., 2001; Weidenbech, Waters, Birget, Broditskly, & Memon, 2005).

The third program concerns the development and evaluation of novel technologies designed to help people understand and manage security and privacy. The work on P3P, the Platform for Privacy Preferences, is an excellent example. Cranor and colleagues have been particularly concerned with the domain of Internet interaction as a site of privacy problems and have developed P3P as a technology to assist end-users in managing their privacy requirements (Ackerman et al., 1999; Cranor & Reagle, 1998). A central focus of P3P and related technologies is the variability in privacy preferences, which, in an e-commerce context, apply both to consumers and to vendors or retailers. P3P provides a mechanism for machine-readable privacy policy specifications, allowing the policies of each party to be compared and potential incompatibilities flagged. A related problem is the issue of personalization in e-commerce; personalization of the e-commerce experience provides value both for the consumer and the retailer but requires collecting and retaining information about browsing and shopping habits that may contradict standard privacy models. Cranor (2003) explores this problem and suggests a range of technological approaches that may mitigate the privacy risks. The recent rise in the number of so-called phishing attacks, in which people are fooled into disclosing personal information at Web sites masquerading as trusted servers, has similarly spawned the development of new security mechanisms (e.g., Dhamija & Tyger, 2005).

Although not part of this emerging interest in usability and security, there are two other bodies of work that are important to note.

The first is that work within the emerging area of ubiquitous computing applications and infrastructures that, although not centrally concerned with privacy and security, has nonetheless needed to tackle these problems as part of broader system development exercises. Examples include infrastructures such as PlaceLab (Schilit et al., 2003), application environments such as ActiveCampus (Griswold et al., 2004), and novel applications such as Reno (Iachello, Smith, Consolvo, Chen, & Abowd, 2005). Our own work has been motivated by participating in research in the area of ubiquitous and pervasive computing.

The second is a broader set of considerations of the relation between information technologies and privacy and most especially the work of those authors who examined the social contexts within which information technology is embedded, including Phil Agre (1994), Roger Clark (1988, 1994), Andrew Clement (1994), and Bruce Schneier (2000). Writing from different perspectives, and with consequences for organizational practice, public policy, social analysis, and information systems design, these authors and others have cast valuable light on the ways in which privacy and security are embedded within broader patterns of social engagement. Our approach here attempts to take this further by, first, elucidating some particular social and cultural practices achieved through the strategic figuring of both privacy and security and, second, by relating this specifically to design practice in information systems.

3. THREE MODELS

Although privacy and security have recently become objects of HCI research attention, as noted previously, the problems remain significant, and the solutions brittle. Following writers such as Agre and Clark, we argue that this is a consequence of too narrow a focus on privacy or on security. This is not simply an argument that privacy and security have a social origin, which has been widely remarked (Flechais et al., 2003; Palen & Dourish, 2003; Weirich & Sasse, 2001); rather, it is an exploration of the nature and scope of this origin. We begin by distinguishing amongst three parallel approaches to thinking about people and privacy, and what we want to think of collectively as information practices.

3.1. Privacy as Economic Rationality

Most discussions of privacy adopt, either explicitly or implicitly, an approach that we refer to as an economic model. Economic here does not mean financial; rather, the model is economic in that its central element is the idea

of a trade-off between risk and reward, the cost and benefit associated with sharing, revealing, or transmitting information. Information is modeled as a commodity that can be traded. Discussions of credit card use, for example, regularly turn on the idea that the benefit that people gain from being able to charge purchases conveniently outweighs the potential costs of making information about purchase history available to the credit card company; similar arguments apply to situations as diverse as store loyalty cards and presence availability in Instant Messaging (Patil & Lai, 2005). We refer to this as an economic model because of its fundamental reliance on two concepts. The first is the concept of exchange value; this model implies both that information is traded for benefit and that items of information can be compared and ordered by their exchange values. The second is the figure of the rational actor, the user who assesses current and potential future situations and calculates their costs and impacts. The economic approach to privacy models collective action as the outcome of individual decision making by rational actors optimizing for individual benefit.

The economic model is at the heart of many proposals for interactive and ubiquitous computing systems. For instance, Place Lab (Schilit et al., 2003) is a sophisticated platform for location-based services. Privacy is a central consideration and, unlike many systems for location-based applications, Place Lab takes pains to avoid the typical Panopticon approach in which a central system component maintains a record of individual locations and movements. Instead, Place Lab allows a device to become aware of its own location through a range of location technologies. A number of the applications that we might want to build on top of such an infrastructure, however, will involve disclosing individual location. Some effort, then, has gone into developing a framework on top of Place Lab that provides end-users with some control over the ways in which their location is disclosed (Hong et al., 2003). For instance, the research team describes the Place Bar, which allows users to control the degree of specificity with which their location is reported—for instance, at the level of a room, building, street, or city. Exchange and ordering are quite explicit here; a trade-off is made between privacy and specificity, and information is ordered in a hierarchy of ambiguity.

A similar approach uses *k*-anonymity as a means to manage information flow in location-based context-aware systems (Gruteser & Grunwald, 2003). In a *k*-anonymous location-based system, when an individual's location is reported, the degree of accuracy of the information is dynamically adjusted depending on other people in proximity. The intent is that a location report is not enough to uniquely locate a single individual; rather, the report identifies a region occupied by k individuals. When k is large, the region reported is relatively large (although smaller in highly populous areas); as k gets smaller, the information is more accurate and specific to an individual.

By choosing an appropriate value of *k,* a user can trade-off the accuracy of the services delivered via this location information against his or her personal privacy preferences.

Floerkemeier, Schneider, and Langheinrich (2004) adopt a similar approach quite explicitly in their consideration of the use of RFID devices in, for example, creating shopping profiles in physical stores; as they explain in their motivation, they take the position that, as consumers, we "engage in meaningful exchanges that conditionally lead us to disclose parts of our personal data to service providers in return for more or less tangible benefits" (p. 215).

The economic approach has an intuitive appeal, and in a number of cases it is clearly an appropriate model. However, as a sole explanation for privacy practices, it has a number of problems, both conceptually and, consequently, as a model for design. Studies of actual practice fail to display the sort of rational trade-off that this model would suggest (Spiekermann, Grossklags, & Berendt, 2001). Recent research in the area of behavioral economics suggests that traditional rational actor approaches fail to adequately account for everyday behavior even within their own terms of reference (Rabin, 1998.) The notion of stable exchange values for goods, services, and labor on which conventional economic modeling is based seems to fare poorly when applied to human actors who are meant to embody these principles. Instead, psychological and social factors seem to interfere with the mathematical principles of neoclassical economics. In a simple example, although you might pay a neighborhood child $20 to mow your lawn, you would be less likely to mow your neighbor's lawn for $20. Recent approaches that attempt to incorporate psychological elements into economics models, such as prospect theory, revise traditional notions of commodity and expected utility (Kahneman & Tversky, 1979).

More problematically, however, we suggest that economic models fail to recognize that privacy is, essentially, a social practice. A trade-off or exchange between rational actors surely fails to capture the sharing of intimate secrets between lovers (Richardson, 1988) or the morality of full disclosure in closed groups (Kleinman & Fine, 1979). Economic models may provide a gloss or explanatory account of information practices, but accounts and motivations must be distinguished (Schutz, 1943). Privacy is not simply a way that information is managed but how social relations are managed. By shifting focus, we want to look at privacy as part of a range of social practice, rather than focusing on the narrow range of activities (e.g., disclosure of credit card information) to which the economic approach is traditionally applied. Although it is possible to incorporate some of these elements into economic discourse (e.g., in discussions of social and cultural "capital"), complementary models can help us to look at privacy as a part of everyday life.

3.2. Privacy as Practical Action

The second approach is to think of security as a practical phenomenon. This turns attention away from abstract information exchanges and toward the practical detail of what people do.

This is not simply a shift from theory to practice but rather, in an ethnomethodological spirit, is an attempt to find the ways in which security is manifest and produced as a property of mundane settings and everyday practice (Garfinkel, 1967). Security, by this argument, is not an abstract feature of ideal settings; it is a practical, ad hoc, in-the-moment accomplishment of social actors, achieved through their concerted actions in actual settings. Privacy and security are witnessable features of working settings, available at-a-glance as situated practices of looking-and-seeing. When we take this perspective, a quite different set of questions emerge. How do people go about doing work securely? How is the difference between "public" and "private" demonstrated in the ways in which they go about their business? How are private matters organized and accountably produced?

The focus on practice has two major implications. The first is that privacy and security are continual, ongoing accomplishments; they are constantly being produced and reproduced. This is a significant departure from technical models that suggest that your security or privacy needs can be "set up" through a control panel and then left alone; instead, it posits privacy and security as ongoing features of activity, which must always be done securely or in ways that are accountably private and so forth. The second is that they are pervasive elements of everyday settings, which extend beyond the boundaries of any or all computer systems and incorporate organizational arrangements and practices, the physical environment, and so on. Empirical investigations into everyday encounters with information security have documented a number of practical methods by which the need to be able to work in ways that are accountably secure can be achieved. These might involve the delegation of responsibility to other elements of the setting (including technology, people, organizations, and institutions), through the reconfiguration of the content of the activity itself (through the development of conventions or procedures for dealing with partial information), through transformations of the working context (such as the intersection between electronic and physical spaces), and so forth (Dourish, Grinter, Delgado de la Flor, & Joseph, 2004).

3.3. Privacy as Discursive Practice

A third approach to privacy and security is as a discursive phenomenon. Language does not simply describe the world but is part of the process of constituting and shaping the world that we experience. So, the issue here is to understand how the notion of privacy and security are used to categorize ac-

tivities, events, and settings, separating acceptable (secure) actions from unacceptable (insecure) ones.

Clearly, any such use of language embodies a particular perspective; security of what, for whom, and from what? What risks are implied? And who gets to define them? For instance, much discussion of information security occurs in corporate contexts, and corporate security directives typically place organizational conveniences ahead of personal ones. Adams and Sasse (1999) report on a particularly telling case. At one of their sites, people reported sharing their personal passwords with coworkers as part and parcel of being a "team player," to get access to shared data and to get their work done more effectively. The system's users perceived organizational policy concerning personal passwords as being incompatible with their working practice. In this case, then, two models of security are in conflict. To the organization, the primary concern is with information security and the dangers of the accidental disclosure of customer information; to the workgroup, the primary concern is with job security and the dangers of inadequate work performance. Because these models of security are in conflict, only one can be formally recognized; the organization's model is formally sanctioned, even as the workgroup's model is practiced. In this case, the determination of what information security means is one that not only condones some behaviors while outlawing others but also reflects power differentials within the organization.

3.4. Reframing Privacy and Security

Why explore these views? The three perspectives complement each other and reveal different aspects of the settings within which privacy and security are managed. Largely, however, it is the first view that has held sway in information systems design; mechanisms such as access control, for example, are designed as encodings of exchange value. The latter two alternative approaches share a focus on information practices as collective rather than individual phenomena and so turn our attention in different directions. In particular, this third approach to privacy and security as a discursive phenomenon places primary emphasis on the broader social and cultural logics of security—the contexts that shape the distinctions between secure and insecure. When looking at these contexts, it becomes clear that privacy and security cannot be analyzed independently but must be considered alongside such related concerns as risk, danger, secrecy, trust, morality, power, identity, and so on.

4. RISK, DANGER, AND MORALITY

Although there is an extensive literature on risk, and in particular on social theories of risk, this rarely features as part of the discussion of privacy and se-

curity technologies. Any notion of security, however, must, implicitly or explicitly, turn on questions of risk and danger, in particular, those risks against which we must be secure. To formulate a state as "secure" is to define what dangers are faced. Similarly, defining particular states or activities as "insecure" marks them as inappropriate and sanctionable. The example provided by Adams and Sasse (1999) demonstrates the way in which technical notions of security are manifestations of (possibly competing) social accounts of danger.

What this suggests, then, is that the ways in which risks are formulated and used to define states of security are ways in which socially acceptable and unacceptable behavior is demarcated. It is valuable, then, to pay attention to the sources of these attributions.

Of particular interest here, security is defined with respect to a set of perceived risks. Douglas and Wildavsky (1982) explore the cultural aspects of risk formulation and risk selection. They make two primary observations. First, they note that the selection of particular activities and objects as "risky" and matters of concern, and the passing over of other activities and objects as not worthy of being labeled risky is not a purely rational or objective process but rather reflects cultural, political, and moral judgments. Shapin's (2004) recent comments about the morality of diet are a nice case in point; pointing to the peculiar symbolism of food, he notes a transformation in how the risks of diet have been discussed, from a 19th century (and before) model in which obesity was the morally reprehensible consequence of gluttony and moderation was to be exalted, to present-day diet books that offer the ability to "eat as much as you like," moving the responsibility for obesity from willpower of a person to metabolism of a body or to foodstuffs themselves (e.g., "bad carbs"). The source of risk has been transformed, and, in so doing, the risk has been changed from a moral and spiritual one to a chemical and physiological one, in line with changing cultural attitudes. This is not an isolated example, we can see this all around us. Consider the role of American cultural attitudes toward individual mobility, justice, and technological progress in debates around the risks of transportation, the death penalty, or nuclear power.

The second major issue that Douglas and Wildavsky (1982) explore is the way in which risk perception should be read relative to social structure. If we read risks as potentially endangering not individuals but rather social structures and "cultural truths," their second observation takes this further; they note that different social collectives will have different interpretations of risk depending on their position relative to the social structures that might be in question. They spend some time exploring alternative perceptions of risk by those who are placed in more or less central and stable social positions. This is reminiscent of Brian Wynne's (1992) discussion of scientists' and farmers' relative knowledge of nuclear technologies and agricultural practices in discus-

sions about the impact of the Chernobyl disaster. Wynne's studies point to the different interpretations of risk between those in more central and marginal societal positions, as well as pointing more generally to the tension between "inside" and "outside" knowledge when epistemic communities encounter each other. For these different epistemic communities, the rhetoric of risk may take quite different forms. For example, in their discussion of the rhetoric of risk in debates over finding sites for storage facilities for low-level radioactive waste, Bedsworth, Lowenthal, and Kastenberg (2004) point not only to the way in which those who opposed the development of new facilities pointed to the potential risks to the ecosystem and local residents but also to the way in which those who were more invested in scientific accounts of the safety of the facility pointed to the risks to those who stood to benefit from the activities that might generate the waste, such as cancer patients suffering as a result of the constraints on scientific and medical research into treatment strategies. The debates around risk become, in this case, struggles over the rhetorical strategies by which risks are defined; in turn, these strategies are associated with different epistemic communities.

What is important to note here, then, is the way in which definitions of risk mark distinctions between acceptable and unacceptable behavior and, at the same time, relationships between different social groups. For example, Day (1995) focuses on another aspect of the social context of risk perception in her study of men's attitudes toward women's vulnerability in public space. In her study, the construal of women as being "at risk" in public spaces serves as an opportunity for the performance and construction of a range of forms of masculinity—"badass" masculinity, chivalrous masculinity, and so forth. She notes not just the different perceptions of women's vulnerability but how they act as a site to reinforce and reproduce cultural logics of action and interaction. As she notes elsewhere (Day, 2001), these cultural logics of risk also serve to mark regions of the environment as "off limits" and act as a form of social control for the at-risk group. The very fact of risk assessment, indeed, is a means by which a socially meaningful designation of at risk can be formulated, in much the same way as disease risk factors can create a new social category of the "presymptomatic ill" (Fosket, 2004; Lock, 1998; Parthasarathy, 2004).

In the case of information system security, these same considerations are at work (Dourish et al., 2004; Weirich & Sasse, 2001); this is hardly a surprise, given the basic relation between security and social structure that we are drawing. Researchers in privacy and security recognize, of course, that assessments of risk are highly variable and relative. However, our point here is that such assessments are collective rather than individual phenomena. The particular significance of this observation for technological settings revolves around the enacted aspects of social practice; that is, the social meaning of the

distinctions being drawn is not simply a feature of the natural world but is continually reproduced in everyday social behavior. Because security tends to feature in design deliberations as a natural fact rather than a social accomplishment, alternative formulations of security and privacy, and the emergence of norms of action and interpretation as a consequence of social engagement, are typically erased. If we accept that privacy and security are formulated relative to normative conventions of risk and danger and that, in turn, these both give rise to and are shaped in everyday social action, then the design question becomes not how we can reflect social norms within information systems but how we can reconfigure information systems as sites for the production of social and cultural values.

Beck (1992) perhaps suggested the most extensive set of relations between risk and social structure. He argues that essentially the distribution of goods that has traditionally characterized industrial modernism is being displaced or augmented by a distribution of risks, yielding what he describes as the "risk society." This is a jumping-off point for an exploration of transformations of various elements of social life in late modernity, including identity, knowledge, and labor. Beck and others who have adopted his approach place risk at the center of modern social life.

5. SECRECY, TRUST, AND IDENTITY

It is impossible to talk about the keeping and sharing of secrets without talking about those groups among whom secrets are shared or from whom secrets are kept. Secrecy, identity, and affiliation are intimately related. Secrets express intimacies and mark groupings, dividing the world into "us" and "them"—friends, families, fraternities, and more. Indeed, the common feature of studies of secrecy is the way in which the practices of keeping and sharing secrets are ways in which affiliation and membership are managed and demonstrated. Again, the use of information to demarcate boundaries is no surprise when we look at the sociological literature; what is important here is how these views can illuminate contemporary computational practice.

There are two levels on which we can explore the relation between secrecy and collective identity. First, on a relatively superficial level, secrecy is intimately connected with social boundaries. Secrets—shared information whose disclosure would somehow endanger the parties to its knowledge—simultaneously cement a bond between those who share it and mark their difference from those with whom it is not to be shared, and this operates across a range of settings from secret societies (Erickson, 1981) to illicit relationships (Richardson, 1988). A second approach is more relevant to our interests here, however, which is to examine the notion of "cultures of secrecy"—social settings that give meaning to certain kinds of information, denoting them as se-

crets and hence giving meaning to patterns of information sharing. We note two aspects here.

Cultures of secrecy make information and its flow meaningful to those who are part of them, and the flow of information itself serves to reproduce those cultures. Merten (1999) introduces this concept in an analysis of information practices in high schools. Unsurprisingly, secrets are used to reinforce and rebel against authority relationships between children and parents, teachers, and adults, but they are also an important resource in managing and navigating the complex world of peer social relations. Secrecy itself, as a marker of a social relationship, is frequently in these cases more important than the content of the information; secrets may be used strategically to cement alliances and deepen friendships. Of particular concern is not just the fact of secret-sharing as a cultural marker of intimacy but the process by which people learn how to share, keep, and use secrets and how the dynamics of peer relations and family relations are sites for the negotiation of norms about what is to be shared and under what circumstances. Part of the process of keeping a secret is recognizing one in the first place, which requires a sensitivity not only to the information itself but to the costs of disclosure (which may themselves lie largely in risks to other social relationships). What is important here is not the secrets themselves but the collective orientation toward practices of secrecy.

Indeed, knowing how to treat information—as sensitive or not—becomes a marker of membership and competent practice. In an ethnographic study of amateur mushroom enthusiasts (Fine & Holyfield, 1996), trust and secrecy both play an important role in the development of group cohesion. As new members join the group, they must learn to trust in others' identification of edible and inedible mushrooms and their use of them in various dishes produced for general consumption, and, at the same time, they must also learn the group's conventions and practices toward members' knowledge of particular mushroom-collecting spots (which are highly personal and carefully guarded). Asking someone for their favorite spots (or rather, expecting to be told someone's favorite spots) is highly inappropriate; at times, members will go to lengths to avoid being heard as asking for this information. As part of the process of enculturation, new members must learn what sort of information is to be shared and what is not to be and must develop new understandings of the norms that govern information use.

In the area of information systems, our own studies, reported in more detail elsewhere, suggest that this may be a fruitful line of inquiry (Anderson & Dourish, 2005). In a study that we conducted of long-haul truckers, similar practices of trust and secrecy proved important in the social life of the group. As people become truckers, they learn from other truckers the secrets of running with heavy loads, driving longer hours, and navigating effectively and

cheaply with wide loads, as well as culturally appropriate control over information flows, such as temporarily disabling GPS cab monitors or respecting conventions about known secrets. For example, asking a trucker for a contact in a city to get a return load, or for a password to a Web site that provides return loads, is extremely inappropriate. In truck stops where a trucker would be using a laptop, other truckers coming into the space would go to great lengths to avoid looking at the screen and would probe to be sure that the trucker was not working to find a load in case they seem to be violating this convention. Truckers' collective, normative information practices define and mark group membership.

These practices may also manifest themselves in the practice of "not noticing" information in the first place, as in a study we conducted of seniors in an assisted care facility (Anderson & Dourish, 2005). The facility had attached load sensors to the seniors' beds to track their weight for health reasons. However, the information so gathered and reported to family members could also indicate that their parent was sleeping with someone (and, through the use of RFID, could potentially specify who). Family members were upset with the management for telling them this information, so it was not reported again; it became invisible information. This was in essence both a collective decision and one that reflected the power dynamics at work (because the family members pay the bills and so hold ultimate power). In fact, this constituted the reemergence of previously agreed-on information practice of the old community. The service people on hand in the house had always known this type of information but chose "not to see it" unless they were explicitly asked. They felt it was not their job to do so. It was agreed-on invisible information, a "safe" behavior in that context.

Appropriate information system design, then, recognizes that information practices—the selective sharing of information and its appropriate management, including forgetting and not noticing—not only are embedded within social groups but are ways that the distinctiveness and boundaries of groups may be identified and reinforced. Appropriate hiding and sharing of information is a marker of social affiliation and a way that membership is accountably demonstrated. Objects and artifacts are not inherently public or private; rather, these categories are negotiated in use as information is strategically deployed to shore up or break down boundaries between people and social groups (Nippert-Eng & Melican, 2005). Information technologies provide new ways to turn identity into an actively managed component of social life; the use of multiple SIM cards in mobile phones, for example, allows individuals to carefully manage their accessibility at different times and in different places (Green, 2002); Instant Messaging, similarly, provides new ways to manage presence and to negotiate participation in different social groups (Grinter & Palen, 2002). The issues of identity work in information practices and the assumptions behind technological designs are perhaps most strik-

ingly illustrated by looking at other cultures (Bell, in press). The ideas of privacy inscribed in our technologies are derived largely from a Western context in which the individual human is the natural unit of social activity and analysis. However, in cultures in which the family, household, or lineage group play more significant roles, the boundaries across which information flows are radically different. These differences underscore the central observation that information technology does not simply encode social practice but is a site at which it is developed.

6. COLLECTIVE INFORMATION PRACTICE IN DESIGN

As we stated at the outset, topics such as privacy and security are common features of discussions about the design and impacts of new information technologies, particularly in ubiquitous computing settings. Talking about the need to maintain privacy and provide security, however, frames these concepts as stable and uniform features of the world, independent of the particular social and cultural circumstances in which individuals find themselves at particular moments. As we have tried to make clear through this discussion, we need to look in more detail not at privacy and security as absolutes but, rather, at what is being done through those concepts.

It is for this reason that we have been speaking not simply about privacy and security but more broadly about information practice. Practice, in the words of Etienne Wenger (1998), is "first and foremost a process by which we can experience the world and our encounters with it as meaningful" (p. 51). So, practice makes the world meaningful through the ways in which we encounter it as offering particular structures and opportunities, and these are collective experiences of meaningfulness. By information practice, then, we are referring to the ways in which we collectively share, withhold, and manage information; how we interpret such acts of sharing, withholding, and managing; and how we strategically deploy them as part and parcel of everyday social interaction. Our intent in focusing on collective information practice is two-fold. One goal is to contribute to the development of an analytic perspective that sets investigations of technological issues in privacy and security within a broader social context. The second is to suggest that collective information practice might, itself, be a target for design.

Ethnographic and related methods have long been adopted in HCI in pursuit of practice-oriented design (e.g., Anderson, 1994; Blomberg, Suchman, & Trigg, 1997; Hughes et al., 1993; Simonsen & Kensing, 1997), although generalizable methods have proven difficult to formulate. In the approach that we have adopted, we see specific configurations of practice not as targets for design but as the outcome of processes and mechanisms whereby collective practice is developed, shared, sustained, and transformed. For this particular project, that means seeing privacy and security as social products rather

than natural facts, which in turn frames technology as a site at which social meaning is produced. Accordingly, we see the primary design goal as providing people with the resources that they need to enact information practices.

There are many forms that this might take, and our goal here is not to suggest any simple transformation of analytic perspectives into design recommendations. However, we hope to make our ideas clearer by showing how they are currently being put to work in our own design endeavors.

Our current approach is to use visual techniques to provide resources for enacting information practice. Technologically, we draw on design elements such as "accounts" (Dourish, 1997), social translucence (Erickson et al., 1999), "social navigation" (Höök, Munro, & Benyon, 2002), and "seamful" design (Chalmers, 2004). The central design consideration is to make a range of mechanisms and actions visible to render them a resource for action. These include relevant system states, mechanisms, representations of the actions of others, and reflexive accounts of one's own action—features that interface abstractions often hide but that are key to crafting appropriate self-presentations and developing collective practice.

In our current development efforts, we have focused in particular on two design principles—visualizing system activity and integrating configuration and action. Visualizing system activity gives users a means of understanding and assessing the consequences of their action. By providing dynamic feedback on relevant but hidden aspects of system activity, such as network traffic and configurations, we provide people with a means to understand the relation between their actions and the technology configuration through which they are performed. Integrating configuration and action reflects the concentration in our account of information practices on their performance, not their expression or articulation through configuration panels (indeed, such articulation, when it arises, is itself performance). In the area of information sharing, it is artificial to separate expression and performance. In our prototypes, this means replacing the traditional separation between "preference panels" and application windows with interaction designs in which configuration and action are achieved together. Examples of this approach include extensions of direct manipulation to couple both sharing configuration and feedback to representations of digital objects, for instance, so that they not only are "handles" for action but also continually display current and past state. A detailed account of the interaction design is not appropriate here but can be found elsewhere (DePaula et al., 2005).

Strongly influenced by the reconceptualization of privacy and security as collective information practices, our current design efforts provide a testbed for moving from supporting practice to supporting its collective shaping and sharing. This testbed is a partial solution and still under development. More generally, however, these designs provide an initial view of the ways in which a broader view of privacy and security can influence the design process.

7. CONCLUSIONS

It is an article of faith in the HCI community—albeit one learned through some hard lessons—that usability cannot be an afterthought in the design of an interactive system. A system cannot be made usable through the simple addition of a user interface, because usability is not limited to the user interface itself; it is a pervasive feature of system design. Privacy and security are similar; support for effective privacy protection cannot be grafted on to a system because it is a pervasive aspect of how that system is designed. In fact, as we have argued here, it is a pervasive aspect of how the system will be used, the context in which it is put to use, the values that it is used to support, the interpretations that others will make of its use, and so forth. Through a broad examination of related literature, we attempt here to illustrate the inevitable social and cultural embeddedness of questions of privacy and security and to draw out the consequences for how we talk about and design information systems.

First, we showed that rational-actor economic models are inadequate as sole explanations of privacy and security practices because they fail to capture other symbolic and social values of those practices. As Sahlins (1972) argues, social action is never purely utilitarian; it is culturally constituted. This is not simply an argument that social factors are elements in the trade-off of costs and benefits but that these are not individual decisions but collective actions, given form and meaning through the ways in which they produce and reproduce cultural and social values.

Second, then, we suggested that information practice—collectively reproduced understandings of the ways in which information should be shared, managed, and withheld—may be more fruitful than traditional conceptions of privacy and security as ways to think about the broader context within which these issues are embedded. Turning our attention away from privacy as an abstract goal and toward information practices as performative helps us see how information is embedded in a wide range of forms of social action. From a design perspective, it calls into question the separation between configuration and action that characterizes most interactive systems for privacy and security management.

Third, we showed that practices are not simply ways in which information is managed but ways in which social actions are achieved. Any adequate account of privacy behaviors, then, must be grounded in an understanding of the specific social and cultural context within which the activity is taking place.

Fourth, we noted that the many different and dynamic social contexts within which people are embedded, as well as the very notion of practice, imply that information needs and uses are continually subject to change, revision, and reinterpretation. One implication of this is that models that require abstract specification (e.g., traditional access control mechanisms and preferences) are inher-

ently limited; again, the separation between configuration and action renders these problematic as a means to enact information practice.

Fifth, and finally, we attempted to show that information practices cannot be separated from the concerns for risk, danger, trust, secrecy, identity, morality, and power that collectively give them meaning. Privacy is not a concept that can be separated from the collective practices through which it is achieved and made operable or from the other elements that are achieved through those same practices. Schneier (2000) notes that "security mechanisms that aren't understood … by everyone don't work" (p. 373). We have attempted to take this one step further, beginning to unpack the ways in which security and privacy, as information practices, are ways in which people collectively understand the world.

The arguments we present here are intended to further an exploration of privacy and security as social products rather than natural facts. Rather than attempting to encode and replace individuals' or groups' information practices (as conventional privacy technologies do), we can seek to support and augment not just social practices but the means by which they evolve. Our recent work on exploring visualization techniques is geared to enabling people to understand and manage information that requires attention in a situated social context. Given our perspective on information practices as thoroughly situated and contingent ways to achieve concerted social action, the primary goal of this work is to help people to understand the consequences of their actions so that they can be managed appropriately. In essence, then, our approach seeks not to transform privacy into a technical property that can be automated but rather to support the human social and cultural practices through which the whole complex of phenomena—privacy, security, risk, danger, secrecy, trust, identity, morality, power, and so forth—are managed and sustained.

NOTES

Acknowledgments. We are grateful to Genevieve Bell, Simon Cole, Rogerio De Paula, Christena Nippert-Eng, David Redmiles, Yvonne Rogers, Lucy Suchman, Jennifer Terry, and three anonymous reviewers for valuable and insightful feedback on earlier drafts.

Support. This work was supported in part by Intel Corporation and by the National Science Foundation under awards 0133749, 0205724, 0326105, 0527729, and 0524033.

Authors' Present Addresses. Paul Dourish, Donald Bren School of Information and Computer Sciences, University of California Irvine, Irvine, CA 92697. E-mail: jpd@ics.uci.edu. Ken Anderson, People and Practices Research, Intel Corporation, 20270 NW Amberglen Court, MS AG1–110, Beaverton, OR 97006. E-mail: ken.anderson@intel.com.

HCI Editorial Record. First manuscript received January 31, 2005. Revisions received August 27, 2005, and January 17, 2006. Accepted by Lucy Suchman. Final manuscript received March 6, 2006. — *Editor*

REFERENCES

Ackerman, M., Cranor, L., & Reagle, J. (1999). Privacy in e-commerce: Examining user scenarios and privacy preferences. *Proceedings of the ACM 1999 Conference on Electronic Commerce* (pp. 1–8). New York: ACM.

Adams, A., & Sasse, A. (1999). Users are not the enemy: Why users compromise security mechanisms and how to take remedial measures. *Communications of the ACM, 42*(12), 40–46.

Agre, P. (1994). Surveillance and capture: Two models of privacy. *The Information Society, 10,* 101–127.

Agre, P., & Rotenberg, M. (Eds.). (1997). *Technology and privacy: The new technological landscape.* Cambridge, MA: MIT Press.

Altman, I. (1977). Privacy regulation: Culturally universal or culturally specific? *Journal of Social Issues, 13*(3), 66–84.

Anderson, R. (1994). Representations and requirements: The value of ethnography in system design. *Human–Computer Interaction, 9,* 151–182.

Anderson, K., & Dourish, P. (2005). Do you know where your mother [trucker] is? *Proceedings of HCI International.*

Beck, U. (1992). *The risk society: Towards a new modernity.* London: Sage.

Bedsworth, L., Lowenthal, M., & Kastenberg, W. (2004). Uncertainty and regulation: The rhetoric of risk in the California low-level radioactive waste debate. *Science, Technology and Human Values, 29,* 406–427.

Bell, G. (2006) Satu Keluarga, Satu Komputer [One home, one computer]: Cultural accounts of ICTs in South and Southeast Asia. *Design Issues, 22*(2), 35–55.

Blomberg, J., Suchman, L., & Trigg, R. (1997). Notes on the work-oriented design project in three voices. In G. Bowker, S. L. Star, W. Turner, & L. Gasser (Eds.), *Beyond the great divide: Socio-technical systems and cooperative work.* Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Boyle, M., & Greenberg, S. (2005). The language of privacy: Learning from video media space analysis and design. *ACM Trans Computer–Human Interaction, 12,* 328–370.

Brin, D. (1999). *The transparent society.* New York: Perseus.

Brodie, C., Karat, C.-M., & Karat, J. (2005). Usable security and privacy: A case study of developing privacy management tools. *Proceedings of the SOUPS 2005 Symposium on Usable Privacy and Security* (pp. 35–42). New York: ACM.

Brostoff, S., & Sasse, A. (2000). Are passfaces more usable than passwords? In S. McDonald, Y. Waern, & G. Cockton (Eds.), *Proceedings of HCI 2000 Conference on People and Computers XIV—Usability or Else!* (pp. 405–424). London: Springer.

Brostoff, S., & Sasse, A. (2001). Safe and sound: A safety-critical design approach to security. *Proceedings of the ACM 2001 Workshop on New Security Paradigms* (pp. 41–50). New York: ACM.

Chalmers, M. (2004). A historical view of context. *Computer-Supported Cooperative Work, 13,* 223–247.

Clark, R. (1988). Information technology and dataveillance. *Communications of the ACM, 31*(5), 498–512.

Clark, R. (1994). The digital persona and its applications to data surveillance. *The Information Society, 10,* 77–92.

Clement, A. (1994). Considering privacy in the development of multi-media communication. *Computer-Supported Cooperative Work, 2,* 67–88.

Consolvo, S., Smith, I., Matthews, T., LaMarca, A., Tabert, J., & Powledge, P. (2005). Location disclosure to social relations: Why, when, & what people want to share. *Proceedings of the CHI 2005 Conference on Human Factors in Computing Systems* (pp. 81–90). New York: ACM.

Cranor, L. (2003). "I didn't buy it for myself": Privacy and e-commerce personalization. *Proceedings of the ACM 2003 Workshop on Privacy in the Electronic Society* (pp. 111–117). New York: ACM.

Cranor, L., & Reagle, J. (1998). Designing a social protocol: Lessons learned from the platform for privacy preferences project. In J. K. MacKie-Mason, & D. Waterman (Eds.), *Telephony, the Internet, and the media.* Mahwah, NJ: Lawrence Erlbaum Associates.

Day, K. (1995). Assault prevention as social control: Women and sexual assault prevention on urban college campuses. *Journal of Environmental Psychology, 15,* 261–281.

Day, K. (2001). Constructing masculinity and women's fear in public space in Irvine, California. *Gender, Place and Culture, 8*(2), 109–127.

DePaula, R., Ding, X., Dourish, P., Nies, K., Pillet, B., Redmiles, D., et al. (2005). In the eye of the beholder: A visualization-based approach to information system security. *International Journal of Human–Computer Studies, 63,* 5–24.

Dhamija, R., & Tyger, J. (2005). The battle against phishing: Dynamic security skins. *Proceedings of the SOUPS 2005 Symposium on Usable Privacy and Security* (pp. 77–88). New York: ACM.

Douglas, M., & Wildavsky, A. (1982). *Risk and culture.* Berkeley: University of California Press.

Dourish, P. (1997). Accounting for system behaviour: Representation, reflection and resourceful action. In M. Kyng & L. Mathiassen (Eds.), *Computers and design in context* (pp. 145–170). Cambridge, MA: MIT Press.

Dourish, P., Grinter, R., Delgado de la Flor, J., & Joseph, M. (2004). Security in the wild: User strategies for managing security as an everyday, practical problem. *Personal and Ubiquitous Computing, 8*(6), 19–30.

Erickson, B. (1981). Secret societies and social structure. *Social Forces, 60,* 188–210.

Erickson, T., Smith, D., Kellogg, W., Laff, M., Richards, J., & Bradner, E. (1999). Socially translucent systems: Social proxies, persistent conversation, and the design of "babble." *Proceedings of the CHI 99 Conference on Human Factors in Computing Systems* (pp. 72–79). New York: ACM.

Fine, G., & Holyfield, L. (1996). Secrecy, trust, and dangerous leisure: Generating group cohesion in voluntary organizations. *Social Psychology Quarterly, 59,* 22–38.

Flechais, I., Sasse, M. A., & Hailes, S. (2003). Bringing security home: A process for developing secure and useable systems. *Proceedings of the ACM/SIGSAC New Security Paradigms Workshop* (pp. 49–57). New York: ACM.

Floerkemeier, C., Schneider, R., & Langheinrich, M. (2004). Scanning with a purpose—Supporting the fair information principles in RFID protocols. *Proceedings of the UCS 2004 Second International Symposium on Ubiquitous Computing Systems* (pp. 214–231). Berlin, Germany: Springer-Verlag.

Fosket, J. (2004). Constructing "high-risk women:" The development and standardization of a breast cancer risk assessment tool. *Science, Technology, and Human Values, 29,* 291–313.

Garfinkel, H. (1967). *Studies in ethnomethodology.* Cambridge, UK: Polity.

Gavison, R. (1980). Privacy and the limits of the law. *Yale Law Journal, 89,* 421–471.

Good, N., & Krekelberger, A. (2003). Usability and privacy: A study of Kazaa P2P file-sharing. *Proceedings of the CHI 2003 Conference on Human Factors in Computing Systems.* New York: ACM.

Green, N. (2002). On the move: Technology, mobility, and the mediation of time and space. *The Information Society, 18,* 281–292.

Grinter, R., & Palen, L. (2002). Instant messaging in teen life. *Proceedings of the CSCW 2002 Conference on Computer-Supported Cooperative Work.* New York: ACM.

Griswold, W., Shanahan, P., Brown, S., Boyer, R., Ratto, M., Shapiro, B., et al. (2004). ActiveCampus: Experiments in community-oriented ubiquitous computing. *IEEE Computer, 37*(10), 73–81.

Grudin, J. (2001). Desituating action: Digital representation of context. *Human–Computer Interaction, 16,* 269–286.

Gruteser, M., & Grunwald, D. (2003). Anonymous usage of location-based services through spatial and temporal cloaking. *Proceedings of the Mobisys 2003 ACM/ USENIX International Conference on Mobile Systems, Applications and Service.* Berkeley, CA: USENIX

Hong, J., Boriello, G., Landay, J., McDonald, D., Schilit, B., & Tygar, J. (2003). Privacy and security in the location-enhanced World Wide Web. (Seattle, WA). Lecture Notes in Computer Science (LNCS2864). *Proceedings of the International Conference on Ubquitous Computing Ubicomp 2003.* Berlin, Germany: Springer.

Höök, K., Munro, A., & Benyon, D. (Eds.). (2002). *Designing information spaces: The social navigation approach.* London: Springer.

Hughes, J., Randall, D., & Shapiro, D. (1993). From ethnographic record to systems design: Some experiences from the field. *Computer-Supported Cooperative Work, 1,* 123–141.

Iachello, G., Smith, I., Consolvo, S., Chen, M., & Abowd, G. (2005). Developing privacy guidelines for social location disclosure applications and services. *Proceedings of the SOUPS 2005 Symposium on Usable Privacy and Security* (pp. 65–76). New York: ACM.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–292.

Kleinman, S., & Fine, G. (1979). Rhetorics and action in moral organizations: Social control of Little Leaguers and ministry students. *Urban Life, 8,* 275–294.

Langheinrich, M. (2002). A privacy awareness system for ubiquitous computing environments. *Proceedings of the Ubicomp 2002.*

Lock, M. (1998). Breast cancer: Reading the omens. *Anthropology Today, 14,* 8–16.

Merten, D. (1999). Enculturation into secrecy among junior high school girls. *Journal of Contemporary Ethnography, 28,* 107–137.

Millett, L., Friedman, B., & Felten, E. (2001). Cookies and Web browser design: Toward realizing informed consent online. *Proceedings of the CHI 2001 Conference on Human Factors in Computing Systems* (pp. 46–52). New York: ACM.

Nippert-Eng, C., & Melican, J. (2005). Concealment and disclosure: Wallets, purses, and identity work in modern societies. Manuscript in preparation.

Palen, L., & Dourish, P. (2003). Unpacking "privacy" for a networked world. *Proceedings of the CHI 2003 Conference on Human Factors in Computing Systems.* New York: ACM.

Parthasarathy, S. (2004). Regulating risk: Defining genetic privacy in the United States and Britain. *Science, Technology, and Human Values, 29,* 332–352.

Patil, S., & Lai, J. (2005). Who gets to know what when: Configuring privacy preferences in an awareness application. *Proceedings of the CHI 2005 Conference on Human Factors in Computing Systems.* New York: ACM.

Patrick, A., Long, C., & Flinn, S. (2003). HCI and security systems (workshop description). *Adjunct Proceedings of CHI 2003.* New York: ACM.

Rabin, M. (1998). Psychology and economics. *Journal Of Economic Literature, 36,* 11–46.

Richardson, L. (1988). Secrecy and status: The social construction of forbidden relationships. *American Sociological Review, 53,* 209–219.

Riegelsberger, J., Sasse, M. A., & McCarthy, J. (2003). Shiny happy people building trust? P*roceedings of the ACM Conference of Human Factors in Computing Systems CHI 2003* (pp. 121–128). New York: ACM.

Sahlins, M. (1972). *Culture and practical reason.* Chicago: University of Chicago Press.

Schneier, B. (2000). *Secrets and lies: Digital security in a networked world.* New York: Wiley.

Schilit, B., LaMarca, A., Borriello, G., Griswold, W., McDonald, D., Lazowska, E., et al. (2003). Challenge: Ubiquitous location-aware computing and the "Place Lab" initiative. *Proceedings of the ACM International Workshop on Wireless Mobile Applications and Services on WLAN.* New York: ACM.

Schutz, A. (1943). The problem of rationality in the social world. *Economica, 10*(38), 130–149.

Shapin, S. (2004). The great neurotic art. *London Review of Books,* 26(15).

Simonsen, J., & Kensing, F. (1997). Using ethnography in contextual design. *Communications of the ACM, 40*(7), 82–88.

Spiekermann, S., Grossklags, J., & Berendt, B. (2001). E-privacy in second generation e-commerce: Privacy preferences versus actual behavior. *Proceedings of the EC 2001 Conference on Electronic Commerce* (pp. 38–47). New York: ACM.

Want, R. (2004, January). RFID: A key to automating everything. *Scientific American, 290,* 56–65.

Weidenbeck, S., Waters, J., Birget, J.-C., Broditskly, A., & Memon, N. (2005). Authentication using graphical passwords: Effects of tolerance and image choice. *Proceedings of the SOUPS 2005 Symposium on Usable Privacy and Security* (pp. 1–12). New York: ACM.

Weirich, D., & Sasse, A. (2001). Pretty good persuasion: Steps towards effective password security in the real world. *Proceedings of the ACM xxxx Workshop on New Security Paradigms Workshop* (pp. 137–143). New York: ACM.

Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity.* Cambridge University Press.

Westin, A. (1968). *Privacy and freedom.* New York: Atheneum.

Whitten, A., & Tygar, D. (1999). Why Johnny can't encrypt: A usability evaluation of PGP 5.0. *Proceedings of the Ninth USENIX Security Symposium.*

Wynne, B. (1992). Misunderstood misunderstandings: Social identities and public uptake of science. *Public Understanding of Science, 1,* 281–304.