ISR Institute for Software Research
University of California, Irvine

Privacy, Security... and Risk and Danger and Secrecy and Trust and Morality and Identity and Power: Understanding Collective Information Practices

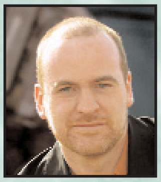

Paul Dourish
University of California, Irvine
jpd@ics.uci.edu

Ken Anderson
Intel Corporation
ken.anderson@intel.com

January 2005

ISR Technical Report # UCI-ISR-05-1

Institute for Software Research
ICS2 210
University of California, Irvine
Irvine, CA 92697-3425
www.isr.uci.edu

http://www.isr.uci.edu/tech-reports.html

Privacy, Security… and Risk and Danger and Secrecy and Trust and Morality and Identity and Power: Understanding Collective Information Practices

Paul Dourish
Institute for Software Research
University of California, Irvine
Irvine, CA 92697-3425
jpd@ics.uci.edu

Ken Anderson
People and Practices Research
Intel Corporation
Hillsboro, OR
ken.anderson@intel.com

ISR Technical Report # UCI-ISR-05-1

January 2005

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 paper, we argue that we can understand privacy and security only by looking at the broader social and cultural contexts within which it is embedded. Privacy and security are difficult concepts to grapple with precisely because they are caught up in larger collective rhetorics and practices of risk, danger, secrecy, trust, morality, identity and more. When we try to separate these concepts, though, the results are incoherent. We argue for a move away from narrow views of privacy and security, and towards situated and collective information practices.

Privacy, Security… and Risk and Danger and Secrecy and Trust and Morality and Identity and Power: Understanding Collective Information Practices

Paul Dourish
Institute for Software Research
University of California Irvine
jpd@ics.uci.edu

Ken Anderson
People and Practices Research
Intel Corporation
ken.anderson@intel.com

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 paper, we argue that we can understand privacy and security only by looking at the broader social and cultural contexts within which it is embedded. Privacy and security are difficult concepts to grapple with precisely because they are caught up in larger collective rhetorics and practices of risk, danger, secrecy, trust, morality, identity and more. When we try to separate these concepts, though, the results are incoherent. We argue for a move away from narrow views of privacy and security, and towards situated and collective information practices.

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 not just increasingly ubiquitous, but increasingly essential for the conduct of everyday life. The inexorable drive of Moore's Law has resulted in computational devices which are ever more powerful computationally and ever more accessible economically, and devices from home appliances to hotel door-knobs increasingly exhibit digital capabilities.

However, although Moore's Law is typically stated in quantitative terms, it has arguably wrought a qualitative transformation in our experience of computation. Computer systems are not simply smaller, faster, and cheaper; they are radically different *kinds* of objects than they used to be. Our expectations for what computers are, and for what we might be able to do with them, have changed fundamentally since the first commercial computers became available. Personal computers no longer primarily "crunch numbers," and the idea of advertising that a computer can "help balance the family accounts" is, these days, endearingly quaint.

Similarly, our society's ability to store, manage, and process information has also changed radically, operating on so different a scale than was previously possible as to open up radical new possibilities that could not previously have been imagined. As storage capacities rise, and as more powerful computational engines make it possible to search, process and filter huge volumes of information, 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 like Walmart can insist that their suppliers include them in all products, easing inventory management but also enabling new forms of electronic monitoring and profiling.

There is a pervasive sense in all this that our new digital capabilities outstrip our understanding of how to use them. One of the major focuses of debates about rapid technology development is the problem of privacy. Ubiquitous computing, with its emphasis on capturing aspects of personal and collective "context" or operating in a pervasive sensing environment, is perhaps one area where these issues are particularly visible, and in this area, privacy is widely acknowledged as a major research issue (e.g. Langheinrich, 2002; Palen and 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 pervasive in their structure and usage models. Despite these developments, though, 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 most explorations of the problems of privacy in information system have attempted to understand privacy as a technological problem but not as a social or cultural problem. Our goal here is not to present a design framework or an empirical account of privacy (although we will touch on teach of these topics), but rather to step back and critically examine contemporary discussions of privacy in interactive systems and ubiquitous computing. By drawing on a range of research literature outside the traditional range of these discussions, we wish to illustrate the broader context within which privacy arises as a feature of everyday life. Understanding this larger context – not just how privacy might be achieved, but what it is used to do – opens up new ways of approaching the problems of privacy in information system design.

Two perspectives are central to this exploration. First, we want to explore the contexts within which "privacy" is situated and the curious manner in which privacy has been separated from a set of inextricably associated issues such as risk, danger, safety, etc. In other words, we want to illuminate the questions of "privacies" by placing them in the broader social and cultural context through which collectives negotiate their meaning.

At the same time, we want to move away from the traditional "social impacts" perspective that characterizes much discussion of privacy, technology and society. The "social impacts" approach focuses on the potential impact of new technologies on social structures and social practices, but in doing so, it takes a narrow view of the relationship between technology and society; taking technology largely as a given, it suggests that the relationship between technology and society is unidirectional. Instead, we adopt a more symmetrical perspective; we are interested not merely in how social structures respond to new technological opportunities, but also how social practice shapes technology and technological discourse, the "social shaping" of technology as well as the "social impacts." Technologies derive their meaning from social and cultural contexts, and it is these contexts that we want to explore.

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 towards technology design, we bring to this investigation the belief that, first, social and technical are irremediably, recursively linked; and, second, that any solution to the privacy conundrum must start with how these two domains have interpenetrated. 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 will draw upon share these cultural roots. Even in localized contexts, cultural practices around privacy are remarkably varied (Altman, 1977).

In brief, the argument we wish to develop is this. Conventional interpretations of privacy and security draw on an economic model to frame privacy in terms of the cost/benefit analysis of information exchange. However, in doing so, they patently fail to account for real-world practice, because the economic model is inherently limited. The issue here is that information has a symbolic as well as an instrumental role, and the cultural practices within which it is embedded give meaning not just to information but to the ways in which it is used. The question is not simply, then, about "privacy" but rather about how people engage in sharing, not sharing, and not noticing information; in *choosing* to share, not to share, or not to notice information; and in *being seen to choose* to share, not to share, or not to notice information. As a means towards a more effective understanding between the relationship between people, information, and technology, we advocate an approach based not on abstract ideas of privacy and security but rather on concrete "information practices." We need to look not only at information flows, but at what ends are achieved through those flows; not simply at privacy, but at how people engage in collective information practice.

HCI Perspectives on Security and Privacy

Recently, a number of researchers have been pioneering a new area for HCI research; the domains of privacy and security. While these are different domains, they share some common properties, such as threat assessments, the predictability (and unpredictability) of needs, the centrality of social interaction, and so forth. Further, they have been traditionally related to each other through their solution strategies; typically, we have attempted to solve privacy problems by importing techniques from the security domain (such as access rights management and cryptographic methods.) A full review of emerging research on the HCI aspects of privacy and security is beyond the scope of this paper, but we shall discuss some of the most significant approaches here.

Whitten and Tygar's (1999) study was the first (and probably most influential) in a series of explorations of the usability considerations in security software systems. Through an experimental evaluation of the use of public key encryption and digital signatures in an email 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.

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 and Sasse, 1999; Weirich and Sasse, 2001); some have reported on empirical investigations of the effectiveness of particular techniques (e.g. Brostoff and Sasse, 2000; Riegelsberger et al., 2003); while others have developed new design models for secure interactive systems (Brostoff and Sasse, 2001; Flechais et al., 2003). An important element of this work, which shall concern us here, is the social embedding of privacy practice; the design models that have emerged from their empirical investigations have repeatedly demonstrated the influence of social and organizational setting on security attitudes and behaviors.

Cranor and colleagues have been particularly concerned with the domain of Internet interaction as a site of privacy problems, and have developed P3P, the Platform for Privacy Preferences, as a technology to assist end users in managing their privacy requirements (Ackerman et al., 1999; Cranor and 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 browing 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.

Three Models

While privacy and security have recently become objects of HCI research attention, as noted above, the problems remain significant, and the solutions brittle. Especially in the domain of ubiquitous computing, privacy and security remain significant problems. Our goal in this paper is to explore why this might be. In particular, we suggest that our traditional approaches to privacy and security are narrow and essentially misconstrue the problem. Focussing on privacy or on security is misleading; we need, instead, to think about the collective information practices that manifest themselves as concerns with privacy and security amongst other phenomena. This is not simply an argument that privacy and security have a social origin (Weirich and Sasse, 2001; Flechais et al., 2003; Palen and Dourish, 2003); rather, it is an exploration of the nature and scope of this origin.

We want to begin by distinguishing amongst three related ways of thinking about people and privacy, and what we want to think of collectively as information practices; one traditional interpretation and two alternative construals.

Privacy as Economic Rationality

Most discussions of privacy adopt, either explicitly or implicitly, an approach that we will refer to as an *economic* model. This is not to imply that it is financial; rather, it is economic in that the central element of this approach 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 and 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.

This 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. However, of course, a number of the applications that we might want to build on top of such an infrastructure will involve disclosing individual location. Some effort, then, has gone into developing a framework on top of Place Lab which 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. The notion of both exchange and ordering is quite explicit here; a trade-off is made between privacy and specificity, and information is organized into a hierarchy of ambiguity.

A similar approach uses *k*-anonymity as a means to manage information flow in location-based context-aware systems (Gruteser and 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*, then, a user can trade-off the accuracy of the services delivered via this location information against their personal privacy preferences.

Or, again, Floerkemeier et al. (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."

The economic approach has an intuitive appeal but, as a sole explanation, it has a number of problems, though, both as a conceptual framework 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 et al., 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 upon 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, while you might pay a neighborhood kid $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 and Tversky, 1979).

More problematically, though, 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 and 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 about how information is managed, but about how social relations are managed. By shifting focus, we need 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. We need alternative models that help us to look at privacy as a part of everyday life.

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 towards 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 concerned 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, etc. 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), etc. (Dourish et al., 2004).

Privacy as Discursive Practice

A third approach is to think of privacy and security as a discursive phenomenon. By this, we mean to draw attention to privacy and security as aspects of communicative and linguistic practice. 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 activities, 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. Weirich and Sasse (2002) report on a case where organizational rules about password privacy conflict with a workgroup’s local practice of sharing passwords in order to get their work done promptly and professionally. What is particularly of interest here is the way that “security” is defined to apply to the organizational goal’s but not to the group’s goals. Or again, Agre (1995) has noted how, in discussing privacy problems in ubiquitous computing systems, designers often frame the problem not as the potential loss of privacy that people might experience, but rather as the potential rejection of the technology that developers might encounter; the problem to be addressed is one of understanding and reassurance, but the technology itself is a given. The discursive practices, and the formulation and logic of the problems reflect and reinforce power relations and conceptual models.What concerns us here is the exercise of concepts such as security, privacy, secrecy, etc., as ways of both making and erasing distinctions. Conflicts between personal privacy and surveillance are one site for these practices, for example. When appeals to the right to privacy are met with the rejoinder that "people who have nothing to hide have nothing to fear," a rhetorical move has been made from questions of privacy to questions of secrecy; while privacy is acknowledged as a right in Western cultures, there is no right to secrecy, and secrecy has a strong moral dimension Warren and Laslett, 1977).

Reframing Privacy and Security

These latter two alternative approaches share a focus on information practices as *collective* rather than *individual* phenomena. However, they 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.

Our argument is not that these are all a single issue, but rather than they are so firmly interconnected, as related cultural practices, that attempts to focus narrowly on one or another in understanding information work will lead to inevitably incomplete accounts of human action and technological requirements. Accordingly, in what follows, we do not attempt to deal with each issue entirely separately (power, for example, is pervasive), but rather, to explore the set of relationships from different perspectives.

Risk and Danger

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 security technologies.

Beck (1992) has perhaps suggested the most extensive set of relationships 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.

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 recent comments (2004) 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 nineteenth 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 which 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 towards 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 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 discussions 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.

Risk and danger inherently involve uncertainty, and so frequently play a role in debates between social collectives whose interests diverge. In these settings, it can be deployed strategically as a part of forming an interpretive frame that fits the needs of different groups. For example, in their discussion of the "rhetoric of risk" in debates over finding sites for storage facilities for low-level radioactive waste, Bedsworth et al. (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.

Day (1995) focuses on another aspect of the social context of risk perception in her study of men's attitudes towards 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, etc. 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.

Relatedly, risk assessments depend on formulations of problems that are inherently partial and reflect the different demands, needs, expectations and assumptions of different collectives and positions. Fosket's (2004) discussion of the development and subsequent adoption of the Gail model for breast cancer risk assessment is an interesting case in point. The Gail model was a breakthrough in breast cancer risk assessment, but based on available evidence at the time and a relatively limited sample set (particularly with respect to demographics). The model determines a numerical probability of breast cancer incidence, and this numerical focus is clearly valuable, making results both portable (Latour, 1989) and comparable (Bowker and Star, 1999). However, what is now the cut-off point for "high risk" (1.7%) was originally a criterion for inclusion in clinical trials, and so subject to a broader range of concerns about applicability across age ranges, etc. The determination of risk, in this case, is strongly tied to a determination of the effectiveness of a particular treatment (tamoxifen.) Not only does risk assessment, here, create a new category, the "pre-symptomatic ill" (Lock, 1998); it also points to an inherent duality of problem and solution. Fujimura (1992) observed the co-construction of problems and solutions in scientific practice; in this case, we note the co-construction of risks and remedies. Similar issues are at work in the use of DNA testing for "genetic predisposition," a case where technology has created, essentially, new risk categories with very different consequences for the different groups who might have to deal with them (Parthasarathy, 2004).

That assessments of risk are highly variable and relative is scarcely a novel observation. More interesting, though, is to look beyond the level of individual self-interest towards risk assessments as socially shared and socially shaped practices. For example, in his study of cigar smokers, DeSantis (2003) describes a set of rhetorical practices by which cigar smoking is framed as being (relatively) safe – for instance, through a distinction between smoking cigars and cigarettes, an acceptance of the inherent riskiness of everyday life, a substitution of threats (smoking for stress), etc. What is critical to his argument is that these are not individual but collective strategies, ways in which the group as a whole collectively asserts and reinforces its identity. Similarly, a study of nightclub doormen (Monaghan, 2002) discusses how their formulations of what sorts of sexual behaviors are "risky" is intimately bound up with the demonstrably, publicly "masculine" aspects of their employment, and with associated images of what it means to be male, and what sorts of behaviors are expected to accompany this. Threats to masculine status may be more salient than threats to individual health (although threats to personal safety at the hands of aggrieved partners play a yet larger role.)

So the issue here is not simply that risk assessments are variable; it is that risk perception and definition is a culturally situated phenomenon. These examples highlight the ways in which designations of risk – of where it occurs, for whom it is a problem, and what the consequences might be – is a means by which authority is demonstrated, legitimacy is established, identity is constructed, and boundaries affirmed.

Privacy and security technologies, then, are amongst a range of strategies deployed to deal with perceived risks, and the perception of risk is an outcome of collective practice. Risk can never be avoided altogether, and in many settings, especially complex ones, casts a permanent shadow over everyday activity. Risk cannot be ignored; some form of mitigation is needed. What is interesting here, then, is to look at the deployment of technologies for privacy and security in the context of broader rituals and practices of risk mitigation. Vaughan (2004), for example, discusses how two high-reliability organizations – NASA and the US National Air Transportation System – deal with risk as an unavoidable aspect of daily practice. Even high reliability organizations, she argues, encounter "routine nonconformability" – "unanticipated events that deviate from organizational expectations" but that occur as "a regular by-product of the characteristics of the system itself." The heightened consciousness of risk in high reliability organization results in operational procedures designed to highlight, address, and limit these conditions; but, ironically, the very fact of these procedures as an everyday, routine occurrence can, in some cases, have the result that the dangerous conditions become normal, acceptable, and tolerable. Essentially, the information and procedures invoked to manage potentially dangerous anomalies take on a symbolic and ritual importance which masks their original, instrumental aim (Feldman and March, 1981), ritual designed to restore order by maintaining and reaffirming the boundary between safety and danger, purity and contamination (Douglas, 1966; Turner, 1969).

Secrecy and Trust

It is impossible to talk about the keeping and sharing of secrets without talking about those groups amongst 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.

In a technical setting, sharing of passwords is perhaps one of the most recognizable issues. Weirich and Sasse (2002) document, amongst other observations, routine sharing of passwords within families and work groups as part of the process of "team work". As in the case of the Cumbrian sheep farmers, the question of power to define the hierarchy of risk is salient here; risks to organizational information and risks to group performance and cohesion.

The use of information to demark boundaries is no surprise when we look at the sociological literature. Secret societies rely on secrecy not only as a means to demarcate between insiders and outsiders but also as a means to strengthen internal bonds (e.g. cell structures in resistance movements) (Erickson, 1981). On a smaller scale, Richardson (1988) documents the critical role of secrets in the creation and maintenance of illicit relationships, in two ways – first, the creation of intimacy through the sharing of individual secrets, and, later, the solidification of a unified identity for the couple through the sharing of the secret of their own relationship and its history. These two aspects of secrecy vividly demonstrate the way that secrets and boundaries are intimately connected.

A study of high school girls catalogs ways in which shared secrets are used to cement and demonstrate social bonds (Merten, 1999). 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 girls 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 towards practices of secrecy.

Similarly, secrecy practices play an important role in the ethnographic study of amateur mushroom enthusiasts presented by Fine and Holyfield (1996). Here, trust and secrecy both play an important role in the life of the group. In particular, Fine and Holyfield point towards the ways in which they are relevant to the development of group cohesion, operating through the tension between these contradictory notions. 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 towards 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 mistaken for 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 not, and must develop new understandings of the norms that govern information use.

In a study that we conducted of long haul truckers, we saw similar practices of trust and secrecy as important in the social life of the group. As people became truckers, they would learn from other truckers the "secrets" of running with heavy loads, driving longer hours, temporarily disabling GPS cab monitors, and navigating effectively and cheaply with wide loads. Collective information practices also dictate what is not to be shared. 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 wasn't working to find a load in case they seem to be violating this convention. Again, truckers have collective, normative information practices that define and mark group membership.

Indeed, secrecy frequently manifests itself in the practice of “not noticing” information in the first place. One example arose in a study we conducted of seniors in an assisted care facility. 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 that it was not reported again; it became invisible information. This was in essence both a collective decision as well as one that reflected the power dynamics at work (since the family members pay the bills, and so hold ultimate power.) In fact, this constituted the re-emergence of previously agreed-upon information practice of the old community. The service people on hand in the house had always know 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 upon invisible information, a “safe” behavior in that context.

This is familiar in Goffman’s (1966) concept of “civil inattention” in public places. In 90 hours of London Underground observation we saw the same thing with people “making out” on the tube. Other riders would take pains *not* to notice them. On a sunny Saturday afternoon on Primrose Hill in London, couples scattered the hillside having “picnics” on blankets. During the course of the afternoon, many ended up making out for large chunks of time. As other park goers passed by they would turn their head not to notice, so this would be a “safe” place for couples to engage in private intimacy in a public space.

Broadly, then, these studies point towards information practices – the selective sharing of information, and its appropriate management, including forgetting and not noticing – as being not only embedded within social groups, but 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.

Morality

Such processes of boundary maintenance often have a strongly moral character. We have already noted the inherently moral component of secrecy. The moral dimensions of information sharing also feature in Kleinman and Fine’s (1979) exploration of “moral organizations” (organizations that, amongst other things, deploy a series of moral rhetorics with the explicit aim of changing the “core self” of recruits – such as health farms and youth organizations.) For instance, they describe how students at a seminar are encouraged to reveal much about themselves as a route towards “self-knowledge”; this happens not only through formal structures such as forms and reports, but also through informal structures such as study groups. Sharing information is drawn into the moral program of the organization. Accordingly, a more traditional selection of what information to share and with whom is cast as an abrogation of a responsibility and a threat to the community as a whole. These threats and moral positions are themselves strong influences on the assessment and management of socially-perceived risk (Cradock, 2004; Warren, 1979).

Identity

Finally, here, a common thread that has run through much of what we have discussed is the way in which the negotiation of boundaries between self and other is a means by which identity is negotiated and marked. Nippert-Eng and Melican (in preparation) explore this “identity work” in the ways in which people manage the boundaries between personal and public information. 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. 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).

The issues of identity work in information practices and the assumptions behind technological designs are perhaps most strikingly 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 where the family, household, or lineage group play more significant roles, the boundaries across which information flows are radically transformed.

From Privacy and Security to Collective Information Practices

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. When we talk about the need to maintain privacy and provide security, however, we talk about 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 practices. 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." 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 practices, 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.

By looking at information practices, we see the ways in which information is incorporated into rituals of everyday life and is managed as a part of everyday ongoing activity. Indeed, the very notion of "information" as something that is formalized and exchanged begins to dissolve when we take this perspective. Information "exchange" or "transfer" is an inappropriate metaphor for cultural practices; the boundary between information and activity often makes little sense in everyday life. Two examples illustrate this point.

The first concerns the use of an innovative location-based computing infrastructure deployed on a university campus (Barkhuus and Dourish, 2004). The system is structured as a platform that supports a range of applications and services, including geo-messaging, support for in-class interaction, and mechanisms to find the location of friends and colleagues. However, in practice, the "buddy finder" functionality was little-used by the undergraduate students who made up the primary target user population. A primary reason for this is that location is simply not problematic for undergraduates in the way in which it might be for the faculty and graduate students who developed the system. Undergraduate students live highly ordered lives which constrain their location on campus at various times during the day; their movements are highly patterned, and they have only limited discretion over their location. They tend, then, to interact in similarly patterned ways; meeting the same people at the same time, eating in the same places and on the same schedules, throughout an academic quarter. Students know where to find each other, because it is part of the nature of being an undergraduate student, part of their institutional relationship to the university, to move in these predicable ways. Accordingly, then, in practice, a friend’s location is not formulated as a piece of information, an unknown to be sought; encounters are accomplished not through “tracking each other down” but rather as a part of institutional life.

The second study (Reddy and Dourish, 2002) reports on a study of information practices in a hospital setting. While it initially set out to study information “seeking,” and so formulated the problem in terms of needs, queries, and responses, the practice that they observed was quite different. Certainly, questions were asked, queries dispatched, and answers provided. However, hospitals are information-rich environments. Charts, graphs, machines, screens, signals, alerts, and read-outs all provide information directly; but so too do the presence and movement of people, the configurations of technology, and the visual appearance of patients. Being able to “read” these cues is part and parcel of competent, diligent work. So, while information “flows” through the hospital, from unit to unit, “information work” could not be separated off from the administration of medical care; it was seamlessly woven into the practical accomplishment of that care.

Giddens’ (1984) Theory of Structuration posits three concurrent relationships between action and social structure – signification, domination, and legitimation. We are not using this full framework here, but we can see these elements at work. Signification is the process by which actions take on meaning; social structures provide an interpretive frame within which actions can be seen as (and demonstrated to be) meaningful. Domination is the process by which a group or individual limits or shapes the agency of another; essentially, it is the collective exercise of power in forcing, requiring, enjoining or encouraging others to act in particular ways. Legitimation is the process by which some courses of action are rendered acceptable (or unacceptable) through their relationship to established, mutually recognizable patterns of action. Metaphorically, if we replace “action” by “information,” we can begin to understand information practices in collectives. By turning analytic attention away from simple economic models of privacy and commoditized information exchange, and towards information practices and the broader ends that they achieve, we have attempted to show these processes at work and the central role that they play in shaping information behavior.

Analysis and Design

As we stated at the outset, our goal here is not to present a series of “design implications” or guidelines for the development of technologies. Rather, it has been to set out an alternative characterization of the problems of privacy and security, and to show how new understandings of the problems of technology, privacy and identity can emerge from this re-framing. However, while we have been focused here on empirical and analytic accounts of social action, it is important too to show that these do have an impact at the level of the systems that we might design. We illustrate this here by showing how these concerns have manifested themselves in an ongoing design activity (DePaula et al., 2005).

In this article, we have suggested that conventional approaches to privacy and security are brittle because they misconstrue the activities in which people are engaged through those facilities. Instead of taking privacy or security as fundamental issues, we have proposed instead a focus on “collective information practices.” The focus of design, then, shifts from how privacy can be achieved to how information practices can be enacted and sustained. These practices are collective, embedded in social and cultural settings, which suggests in turn that they cannot be replaced by but rather supported by computational mechanisms. In our approach, we conceptualize technology as a platform or stage upon which information practices are performed. The primary design goal is to provide people with the resources that they need to enact information practices.

We focus 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 relationship between their actions and the technology configuration through which they are performed. It is important to note that this visualization does not take the form of the sorts of network monitoring that might be employed by system administrators or network managers. Clearly, end users neither understand nor care to understand the details of network operation, and so we cannot assume this level of technical expertise. Nonetheless, we find that people can understand and appreciate the temporal and structural correlations between their activities and the system's behavior. A useful analogy is with driving; even without understanding the detailed operation of the car's engine, transmission, or steerage assembly, a driver can still make use of the sound of the engine, the feel of the clutch, and the feel of the road through the wheel.

Integrating configuration and action reflects the concentration in our account of information practices that they are performed, not expressed (and indeed, that expression, when it arises, is itself performance.) In contrast, the focus on privacy and security as the expression of preferences tends to yield, in current operating system designs, a separation between a control panel where preferences are set, and some separate window or windows within which the activity of the system is performed. This separation is doubly problematic. Not only does it separate two coextensive forms of activity (the act of "sharing" being distributed across the preference window and the system window), but it also separates the expression of preferences from the occasion or situation in which those preferences are to be invoked. Conventional interfaces separate configuration and action in both space and time, although in everyday practice they are one and the same activity. Speaking and vocal modulation (e.g. intonation, volume, etc) are, for example, inseparable aspects of the same activity; similarly, our design approach seeks to make configuration and action part of the same interactional space.

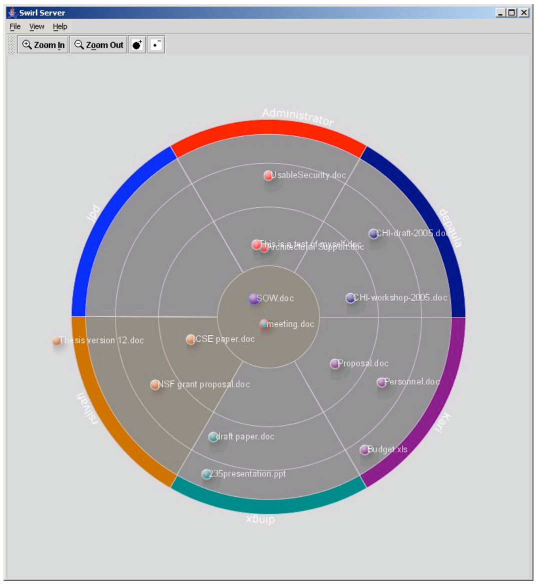

Figure 1: The Impromptu client integrations action, configuration and visualization within a single interface.

We have been applying these principles in Swirl, an infrastructure for usable security, and impromptu, a testbed application. A fuller description of these systems is not immediately relevant here, and can be found elsewhere (DePaula et al., 2005). However, figure 1 shows a screenshot of the impromptu application and its reflection of these principles. Impromptu is a collaborative filesharing application based on a peer-to-peer architecture with automatic host and service discovery. When a set of impromptu clients are on the same network (e.g. when a group of colleagues sits down with laptop, tablet and handheld devices in a meeting room), a shared filespace is created automatically. Inspired in part by Erickson et al's (1999) "social proxy", this filespace is visually represented in the circular region shown in figure 1. Each person sees the same view, to allow for mutually consistent reference (Tatar et al., 1991). People share files by moving them into the circular area; the degree to which information is shared is determined by the proximity of the file icons to the central "shared and persistent" space. So, the visual arrangement of information in this space is simultaneously a view of the system configuration and a way in which actions in the space are carried out. Animation of the space and the objects within it reveals aspects of the system's operation over shared information; while people are not forcibly exposed to the full details of the system's operation, the visual cues are sufficient to illustrate degrees of activity and expected or unexpected events.

Swirl and Impromptu are not finished systems, but prototypes under active development. What is more, they were not designed specifically in response to the approach we have outlined here. However, they have been strongly influenced by the reconceptualization of privacy and security as collective information practices. More importantly here, they also provide an initial view of the ways in which this reconceptualization can influence the design process, and that interactive technologies can be responsive to this sort of critique.

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, and 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, etc. Through a broad examination of related literature, we have been attempting, here, to illustrate the inevitable social and cultural embeddedness of questions of privacy and security, and to draw out its consequences for how we talk about and design information systems.

First, we have shown that rational-actor economic models are inadequate as sole explanations of privacy and security practices because they fail to capture the symbolic and social value of those practices. As Sahlins (1972) argues, social action is never purely utilitarian; culture intervenes. 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 have suggested that "information practices" – 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 towards information practices as performative mechanisms 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 have shown 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 have 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 inherently limited; again, the separation between configuration and action renders these problematic as a means to enact information practice.

Fifth, and finally, we have 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, nor from the other elements that are achieved through those same practices. This becomes, however, an almost intractable problem for technological design. Even if we attempted to design a system to support socially grounded privacy practice, for example, such as system would break down if trafficked in technologically-specified forms of individual or group identity, and so on.

An alternative approach, then, is to reevaluate the relationship between technological and social elements in design. Rather than attempting to encode and replace individuals’ or groups’ information practices (as conventional privacy technologies do), we can seek to support and augment social practices; to provide a setting within which collective practice can operate and evolve.

We hope that we have illustrated how the concepts of information as “private” or “public” do not make sense as abstracted goals. Rather, these are embedded in social, collective and cultural constructs about what information is to be attended to, what are the risks and to whom, and what information is to be ignored. 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, it 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, etc. – are managed and sustained.

Acknowledgements

This work was supported in part by Intel Corporation, and by the National Science Foundation under awards 0133749, 0205724 and 0326105. We are grateful to Genevieve Bell and Yvonne Rogers for valuable feedback on early drafts.

References

Ackerman, M., Cranor, L., and Reagle, J. (1999). Privacy in e-commerce: Examining User Scenarios and Privacy Preferences. Proc. First ACM Conference on Electronic Commerce (Denver, CO), 1-8. New York: ACM.

Adams, A. and 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. (1995). Conceptions of the User in Computer Systems Design. In Thomas (ed), The Social and Interactional Dimensions of Human-Computer Interfaces. Cambridge University Press.

Altman, I. (1977). Privacy Regulation: Culturally Universal or Culturally Specific? Journal of Social Issues, 13(3), 66-84.

Barkhuus, L. and Dourish, P. (2004). Everyday Encounters with Context-Aware Computing in a Campus Environment. Proc. Ubicomp 2004 (Nottingham, UK), 232-249.

Beck, U. (1992). The Risk Society: Towards a New Modernity. London: Sage.

Bedsworth, L., Lowenthal, M., and Kastenberg, W. (2004). Uncertainty and Regulation: The Rhetoric of Risk in the California Low-Level Radioactive Waste Debate. Science, Technology and Human Values, 29(3), 406-427.

Bell, G. (In press). *Satu Keluarga, Satu Komputer,* [One home, one computer]: Cultural Accounts of ICTs in South and Southeast Asia. Design Issues.

Bower, G and Star, L. (1999). Sorting Things Out: Classification and its Consequences. MIT Press.

Brostoff, S. and Sasse, A. (2000). Are Passfaces More Usable Than Passwords? In S. McDonald, Y. Waern & G. Cockton (Eds.), People and Computers XIV - Usability or Else! Proceedings of HCI 2000 (Sunderland, UK), 405-424. London: Springer.

Brostoff, S. and Sasse, A. (2001). Safe and Sound: a safety-critical design approach to security. Proceedings of the ACM New Security Paradigms Workshop, 41-50. New York: ACM.

Cradock, G. (2004). Risk, Morality, and Child Protection: Risk Calculation as Guides to Practice. Science, Technology, and Human Values, 29(3), 314-331.

Cranor, L. (2003). 'I Didn't Buy it for Myself': Privacy and Ecommerce Personalization. Proceedings of the Second ACM Workshop on Privacy in the Electronic Society (Washington, DC.), 111-117. New York: ACM.

Cranor, L. and Reagle, J. (1998). Designing a Social Protocol: Lessons Learned from the Platform for Privacy Preferences Project. In Jeffrey K. MacKie-Mason and David Waterman (Eds.), Telephony, the Internet, and the Media. Mahwah: 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(4), 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.

DeSantis, A. (2003). A Couple of White Guys Sitting Around Smoking: The Collective Rationalization of Cigar Smokers. Journal of Contemporary Ethnography, 32(4), 432-466.

DePaula, R., Ding, X., Dourish, P., Nies, K., Pillet, B., Redmiles, D., Ren, J., Rode, J., and Silva Filho, R. (In preparation.) In the Eye of the Beholder: A Visualization-based Approach to Information System Security. Irvine, CA: University of California, Irvine.

Douglas, M. (1966). Purity and Danger. London: Routledge.

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

Dourish, P., Grinter, R., Delgado de la Flor, J., and 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(1), 188-210.

Erickson, T., Smith, D., Kellogg, W., Laff, M., Richards, J., and Bradner, E. (1999). Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of "Babble." Proc. ACM Conference on Human Factors in Computing Systems CHI'99 (Pittsburgh, PA), 72-79. New York: ACM.

Feldman, M. and March, J. (1981). Information in Organizations as Signal and Symbol. Administrative Science Quarterly, 26(2), 171-186.

Fine, G. and Holyfield, L. (1996). Secrecy, Trust, and Dangerous Leisure: Generating Group Cohesion in Voluntary Organizations. Social Psychology Quarterly, 59(1), 22-38.

Floerkemeier, C., Schneider, R., and Langheinrich, M. (2004). Scanning with a Purpose - Supporting the Fair Information Principles in RFID protocols. Proc. Second International Symposium on Ubiquitous Computing Systems UCS 2004 (Tokyo, Japan).

Fosket, J. (2004). Constructing "High-Risk Women:" The Development and Standardization of a Breast Cancer Risk Assessment Tool. Science, Technology, and Human Values, 29(3), 291-313.

Fujimura, J. (1992). Crafting Science: Standardized Packages, Boundary Objects and "Translation." In Pickering (ed), Science as Practice and Culture. Chicago.

Garfinkel, H. (1967). Studies in Ethnomethodology. Cambridge: Polity.

Giddens, A. (1984). The Constitution of Society. Cambridge: Polity.

Goffman, E. (1966). Behavior in Public Places. New York: The Free Press.

Green, N. (2002). On the Move: Technology, Mobility, and the Mediation of Time and Space. The Information Society, 18, 281-292.

Grudin, J. (2001). Desituating Action: Digital Representation of Context. Human-Computer Interaction, 16(2-4), 269-286.

Gruteser, M. and Grunwald, D. (2003). Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking. Proc. ACM/USENIX International Conference on Mobile Systems, Applications and Services (Mobisys 2003).

Hong, J., Boriello, G., Landay, J., McDonald, D., Schilit, B., and Tygar, J. (2003). Privacy and Security in the Location-Enhanced World Wide Web. Proc. Ubicomp 2003 (Seattle, WA). Springer.

Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision Under Risk. Econometrica, 47(2), 263-292.

Kleinman, S. and Fine, G. (1979). Rhetorics and Action in Moral Organizations: Social Control of Little Leaguers and Ministry Students. Urban Life, 8(3), 275-294.

Langheinrich, M. (2002). A Privacy Awareness System for Ubiquitous Computing Environments. Proc. Ubicomp 2002.

Latour, B. (1989). Science in Action. Cambridge: Harvard University Press.

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(2), 107-137.

Monaghan, L. (2002). Opportunity, Pleasure, and Risk: An Ethnography of Urban Male Heterosexualities. Journal of Contemporary Ethnography, 31(4), 440-477.

Nippert-Eng, C. and Melican, J. (In preparation). Concealment and Disclosure: Wallets, Purses, and Identity Work in Modern Societies. Unpublished manuscript.

Palen, L. and Dourish, P. (2003). Unpacking "Privacy" for a Networked World. Proc. ACM Conf. Human Factors in Computing Systems CHI 2003 (Ft Lauderdale, FL). New York: ACM.

Parthasarathy, S. (2004). Regulating Risk: Defining Genetic Privacy in the United States and Britain. Science, Technology, and Human Values, 29(3), 332-352.

Patil, S. and Lai, J. (2005). Who Gets to Know What When: Configuring Privacy Preferences in an Awareness Application. Proc. ACM Conf. Human Factors in Computing Systems CHI 2005 (Portland, OR). New York: ACM.

Rabin, M. (1998.) Psychology and Economics. Journal Of Economic Literature, 36, 11-46.

Reddy, M. and Dourish, P. (2002). A Finger on the Pulse: Temporal Rhythms and Information Seeking in Medical Work. Proc. ACM Conf. Computer-Supported Cooperative Work CSCW 2004 (New Orleans, LA). New York: ACM.

Richardson, L. (1988). Secrecy and Status: The Social Construction of Forbidden Relationships. American Sociological Review, 53(2), 209-219.

Sahlins, M. (1972). Culture and Practical Reason. University of Chicago Press.

Schilit, B., LaMarca, A., Borriello, G., Griswold, W., McDonald, D., Lazowska, E., Balachandran, A., Hong, J., and Iverson, V. (2003). Challenge: Ubiquitous Location-Aware Computing and the "Place Lab" Initiative. Proc. ACM Intl. Workshop on Wireless Mobile Applications and Services on WLAN (San Diego, CA).

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).

Spiekermann, S., Grossklags, J., and Berendt, B. (2001). E-privacy in Second Generation E-commerce: Privacy Preferences versus Actual Behavior. Proc. ACM Conf. Electronic Commerce EC'01 (Tampa, FL), 38-47. New York: ACM.

Tatar, D., Foster, G., and Bobrow, D. (1991). Designing for Conversation: Lessons from Cognoter. International Journal of Man-Machine Studies, 34(2), 185-209.

Turner, V. (1969). The Ritual Process. Hawthorn, NJ: Aldine de Gruyter.

Vaughan, D. (2004). Organizational Rituals of Risk and Error. In Hutter and Power (eds), Organizational Encounters with Risk. Cambridge University Press.

Warren, C. (1979). The Social Construction of Dangerousness. Urban Life, 8(3), 359-384.

Warren, C. and Laslett, B. (1977). Privacy and Secrecy: A Conceptual Comparison. Journal of Social Issues, 33(3), 43-51.

Wenger, E. (1998). Communities of Practice: Learning, Meaning, and Identity. Cambridge University Press.

Weirich, D. and Sasse, A. (2002). Pretty Good Persuasion: Steps Towards Effective Password Security in the Real World. Proc. ACM New Security Paradigms Workshop, 137-143.

Wynne, B. (1992). Misunderstood Misunderstandings: Social Identities and Public Uptake of Science. Public Understanding of Science, 1(3), 281-304.