Noname manuscript No.
(will be inserted by the editor)

Toward Understanding Distributed Cognition in IT Security Management: The Role of Cues and Norms

David Botta, Kasia Muldner, Kirstie Hawkey, Konstantin Beznosov
{botta,kmuldner,hawkey,beznosov}@ece.ubc.ca
University of British Columbia
Vancouver, Canada

Received: date / Accepted: date

Abstract Information technology security management (ITSM) entails significant challenges, including the distribution of tasks and stakeholders across the organization, the need for security practitioners to cooperate with others, and technological complexity. We investigate the organizational processes in ITSM using qualitative analysis of interviews with ITSM practitioners. To account for the distributed nature of ITSM, we utilized and extended a distributed cognition framework that includes as key aspects the themes of cues and norms. We show how ITSM challenges foster under-use of cues and norms, which comprises a type of risk that may result in outcomes that are adverse to the organization's interests. Throughout, we use scenarios told by our participants to illustrate the various concepts related to cues and norms as well as ITSM breakdowns.

Keywords computer supported cooperative work · cues and norms · distributed cognition · risk · information technology security management · mutual understanding · notifications · transactive memory

1 Introduction

The management of information technology (IT) security is distributed across the organization and requires a great deal of collaborative activity (Goodall et al., 2004a; Kandogan & Haber, 2005; Siegel et al., 2006; Botta et al., 2007; Werlinger et al., 2009). A variety of practitioners, from business managers to IT administrators, participate in the design and daily operations of IT security. IT security management (ITSM) entails significant challenges to practitioners, as they must stay on the cutting edge of technology and adapt to the demands of organizational change (Werlinger et al., 2009; Gagné et al., 2008; Werlinger et al., 2009; Siegel et al., 2006). To meet these challenges, security practitioners rely on each other to develop the "big picture" of what is going on and what to do about it. In other words, their work is characterized by distributed cognition (Salomon, 1993; Hutchins, 1995; Zhang, 1998; Ackerman & Halverson, 2004), which may be defined as *"solving problems by collaboration, where none of the collaborators individually can have a full appreciation of the problem*" (Busby, 2001, p. 238).

Little attention has been devoted to distributed cognition in the context of ITSM in general, and its breakdowns in particular. To address this void, the goals of the research reported in this paper were (1) to provide a greater understanding of organizational processes that influence distributed cognition in ITSM, and (2) to investigate how breakdowns in distributed cognition in ITSM occur.

Our research was qualitative in nature. We relied on a collected interview corpus from the HOT Admin project, which investigated the human, organizational, and technological factors that impact security practitioners in their daily tasks (Hawkey et al., 2008). The analysis reported here is based on data from 35 IT professionals interviewed for the project. Using qualitative description (Sandelowski, 2000), we analyzed our data in terms of Busby's (2001) characterization of distributed cognition, focusing on instances of *cues* and *norms*, two key themes of his framework. With respect to ITSM breakdowns, we identified three rich accounts of security-related incidents in the interview data. We analyzed these breakdowns from a distributed cognition perspective; that is, we focused on cues and norms, and

Address(es) of author(s) should be given how ITSM challenges undermine their effective deployment.

In order to convey distinctions in our findings that are important for understanding ITSM, we employed a modified grounded theory approach to generate subcategories of cues and norms. In particular, we found that cues in ITSM include (1) cues that are not explicitly directed, such as *quick views*, proofs of reliability, and reminders & hints; and (2) notifications that are explicitly directed, such as scripted notifications, notes to self, and escalated notifications. Furthermore, we found that ITSM norms include (1) notification procedures; (2) methods to maintain consistency, such as templates, audits, policies, and standards; (3) establishment of mutual understanding by means of risk assessment, promotion of security awareness, and professional collaboration; and (4) employment of *transactive memory* (Wegner, 1986) to activate the specialized knowledge and skills of others in a group.

This paper makes the following contributions: (1) a confirmation of Busby's focus on cues and norms as central to the functioning of distributed cognition in ITSM; (2) an empirically-grounded elaboration of cues and norms concepts employed in the domain of IT security management, which implies that the two concepts can be useful for analysis of other, similarly complex domains; and (3) a data-driven analysis of how the under-employment of cues and norms leads to the breakdown of distributed cognition, which may lead to increased organizational risks.

In Section 2, we provide the background and related work that motivated our research questions and approach, which are described in Section 3, along with details about our participants. In Section 4, we present our findings (examples of cues and norms, refined to convey important themes in the data); while in Section 5, we provide three illustrations of challenges to ITSM causing under-use of cues and norms that lead to risk and worse. In Section 6, we discuss conclusions drawn from the results and their implications; finally, in Section 7, we discuss the limitations of our approach and the opportunities for future research.

2 Background and Related Work

Traditionally, the focus of IT security research (e.g., (Chebrolua et al., 2005; Fuchs & Pernul, 2007)) has been on devising technical solutions for IT security (e.g., firewalls, intrusion detection systems, and tools for finding vulnerabilities in IT infrastructure) without investigating the factors that influence the usability of these systems. Recently this focus on technical factors was confirmed via a literature survey comparing the relative amounts of IT security research focused on each of technological, human, and social factors (Beznosov & Beznosova, 2007). Some preliminary strides have been made in broadening the scope of IT security research by exploring how organizational and human factors influence security practitioners, the individuals responsible for maintaining security within their organizations. To date, however, little work has been done on understanding the work of security practitioners from a distributed cognition perspective; this is a gap our work aims to fill.

2.1 IT Security Management

We begin by reviewing the work we conducted under the HOT Admin project (Hawkey et al., 2008), which involved a broad study of security practitioners. Specifically, we conducted 34 semi-structured interviews with 35 security practitioners (two of them in one interview) from 16 organizations in 11 industry sectors. We analyzed transcripts of these interviews using qualitative description according to a number of themes central to furthering the research community's understanding with respect to ITSM, including *tasks and tools, security vs. other IT, interactions* and *challenges.* Throughout our analysis, the goal was to answer research questions related to a given theme from a holistic perspective, taking into account the human, organizational, and technological factors at play.

The analysis from the *tasks and tools* theme characterized the workplace of our participants by their responsibilities, goals, tasks, and skills (Botta et al., 2007). The theme of *security vs. other IT* focused on investigating the differences between security and other IT professionals (Gagné et al., 2008); the findings show that security professionals have to contend with a higher level of complexity than other IT personnel. The *interactions* theme (Werlinger et al., 2009) identified nine activities that involve collaboration and cooperation between security practitioners and other stakeholders, such as end users, managers, and other specialists. Analysis of the tools used for these interactions (Werlinger et al., 2009) found that existing tools are not adequate for supporting security practitioners during the various interactions. The *challenges* theme involved an in-depth analysis of the challenges related to the practice of IT security within organizations, including the interplay between human, organizational, and technical factors (Werlinger et al., 2009). As shown in Table 1, the majority of the challenges identified fall into the *organizational* category. Many of the identified challenges stem from the fact that people are often unaware of

Table 1 Challenges for implementing security controls

Human	Lack of training or experience Culture within the organization Communicate security issues
Organizational	Risk estimation Open environments and academic freedom Lack of budget Security as low priority Tight schedules Business relationships with other organizations Distribution of IT responsibilities Controlling access to sensitive data Size of organization Top management support
Technological	Complexity of systems Vulnerabilities (systems/applications) Mobility and distributed access Lack of effective security tools

the knowledge they possess or how it could be valuable to others, something referred to as *tacit knowledge* (Polanyi, 1966). A great deal of ITSM involves tacit knowledge, since articulating or documenting all experiences and/or data is often impractical.

One of our findings is that ITSM is highly distributed in nature (Botta et al., 2007; Werlinger et al., 2009), corroborating prior work (Goodall et al., 2004a; Kandogan & Haber, 2005). In many organizations, departmental units share the IT networks and systems; within each department, one or more individuals are responsible for the local IT infrastructure. The security practitioners we interviewed found this distribution to be a challenge, as it diminished the capability of the organization to apply IT security controls, leading to risk (Werlinger et al., 2009). In the next section, we will see how this particular aspect of ITSM shaped the cognitive framework we relied on for our analysis.

Other researchers have also highlighted the challenging and multi-faceted nature of ITSM. A contextual inquiry with 30 security practitioners at three organizations (Siegel et al., 2006) found that security practitioners face challenges related to the distributed nature of ITSM, lack of access to security training, and the perception of security as an obstacle rather than an enabler of business. Analysis of data from nine semi-structured interviews on intrusion detection work (Goodall et al., 2004a,b) showed that this work is challenging due to the following two factors. First, in addition to highly technical knowledge, security practitioners must also have extensive organizational knowledge about the systems and users within their organizations. Second, ITSM is highly collaborative in nature, requiring security practitioners to interact with a wide variety of stakeholders who have different levels of expertise, and are dispersed throughout the organization. Naturalistic observations of IT administrators (who performed some security duties) in six organizations revealed that they not only need better tool support but they also deal with larger, more complex systems, and face a higher risk of failure than end users (Haber & Bailey, 2007).

Given the challenging nature of ITSM, it is not surprising that security incidents arise (Kraemer & Carayon, 2007; Schultz, 2007). Kraemer & Carayon (2007) show that organizational traits are correlated with certain types of IT security errors. This work focused on overt slips, lapses, mistakes, and violations at the level of individual practitioners and other stakeholders. As we discuss later, we postulate that the challenges to ITSM that make distributed cognition necessary also foster breakdowns of distributed cognition that may result in situations of risk, even when individual practitioners have made no slips, lapses, mistakes, oversights, or violations.

2.2 A distributed cognition framework

As we stated above, one of our key goals was to study the ITSM processes from a cognitive perspective. To do so, given the distributed nature of ITSM, we adopted Busby's (2001) formal framework of distributed cognition. Busby explains that distributed cognition is:

> *concerned with solving problems by collaboration, where none of the collaborators individually can have a full appreciation of the problem. The collaborators can be tools or artifacts of some kind, as well as human information processors, and activity is dynamically referred to parts of the system in both planned and emergent ways. Typically the participants inter-communicate without being fully aware of the extent to which they need to in order to maintain smooth operation of the system* (p. 238).

Busby (2001) relied on a distributed cognition framework to analyze data from 22 interviews with professionals who designed complex processing plants, with the goal of studying how breakdowns lead to errors. Although Busby's study was in the area of design, it is pertinent to the analysis of ITSM for several reasons. First, Busby's scenario of cooperating, specialized units is comparable to the distribution of ITSM in many organizations. In ITSM, individuals from various parts of an organization communicate with each other both formally and on an ad hoc basis to address IT security issues. Second, problem solving in both disciplines shares similar characteristics. ITSM problem solving is characterized by pattern recognition, hypothesis generation and testing, and *bricolage* in situations where success is difficult to ascertain (Botta et al., 2007). These are also characteristics of design problem solving (Simon, 1973; Chandrasekaran, 1990; Goel & Pirolli, 1992). Furthermore, ITSM may involve design responsibilities (Botta et al., 2007).

Busby's (2001) analysis revealed that distributed cognition involves the following two types of phenomena: (1) the occurrence of *cues* (i.e., signals or clues, which participants use to determine when to act and how to act) and (2) the use of *norms* (i.e., standards or patterns regarded as typical, which help make participants' subtasks consistent with each other).

Concerning *cues*, Busby (2001) notes that the nature of cues in conjunction with individuals' prior experiences play a key role in helping individuals interpret what is needed of them. Moreover, differences between individuals limit their ability to provide cues to one another. Busby found that in order to function effectively in a distributed environment, participants need to be aware of the cues that they should provide and of the fact that their own experience may not help them predict what these cues are.

Concerning *norms*, Busby (2001) notes that these are especially important when participants find it hard to predict the effect of their actions on others. He cautions that simply briefing people better, constraining them more, and making norms more prescriptive will not necessarily improve the use of norms, because individuals have difficulty remembering them. Furthermore, Busby emphasizes that people will readily violate conventions if they believe that the norms are unnecessary or if they are working under pressure. This fundamental tension between convention and practical action has been previously discussed in the literature (Suchman, 1983; Poole et al., 1985). Busby argues that the only robust approach is to help people to understand how much they already use norms and to foster sensitivity to their misapplication.

Busby (2001) found both cues and norms to be instrumental in supporting work in distributed environments. Furthermore, cues and norms interact with one another to influence behavior. For instance, a cue may inspire a participant to inspect a norm's underlying assumptions and thereby discover that the norm is outdated. Moreover, when cues are missing or unreliable, norms help participants proceed; and some cues are needed to avoid norm-related errors.

Later work by Busby & Hibberd (2006) helps clarify how the theme of cues and norms contributes to the analysis of distributed cognition. Busby & Hibberd define distributed cognition in terms of Hutchins' (1995) central concern with how information is represented and how the representations are transformed and propagated in the performance of tasks. Busby & Hibberd (2006) point out that the distributed cognition assumption entails that representations are often external (i.e., they lie outside people's heads). They adopt the stance that the main goal of a distributed cognition analysis is to account for how distributed structures are coordinated, and they focus on the role of *organizational artifacts* (e.g., the external norms, rules, schemas, and scripts) to provide coordination. They consider organizational artifacts to be more than determinants of action; rather, actors use them as resources for action, that is, they are subject to an actor's better judgment (see Suchman (1983)). Similarly, Busby & Hibberd (2006) treat organizational artifacts as physical artifacts (e.g. tools), because they are all products of human agency and have a mediating function. In this paper, we include rules, schemas, and scripts as types of norms. Because cues require norms in order to be interpreted, it is important to include them when taking a coordination approach to the analysis of distributed cognition.

There is little prior research on ITSM from the distributed cognition perspective. One observation of security practitioners during a trouble-shooting task (Maglio et al., 2003) describes the failure and repair of the "propagation of representational state" (Hutchins, 1995) due to the failure and repair of mutual understanding (Clark, 1996) between key participants. The repair commenced when *cues* indicated to the participants that they did not have the same mental model. However, Maglio et al. (2003) studied distributed cognition in IT, but did not explicitly use the theme of cues, nor did they generalize mutual understanding to be a kind of norm. Also, they did not explicitly declare the concepts of distributed cognition to be fundamental to their approach. Our study, by explicitly using a distributed cognition framework in the study of ITSM, builds on and extends the work of Maglio et al. (2003).

2.3 Characterization of risk versus error

While we are interested in studying distributed cognition in ITSM, we are also interested in how breakdowns in distributed cognition foster situations of risk and/or error. Since we adopt Busby's (2001) framework, we assume that ITSM breakdowns relate to under-use of cues and norms, as we illustrate later. Situations of risk mean that the security of an organization is in jeopardy even when security practitioners have made no slips, lapses, mistakes, oversights, or violations. While situations of risk may result in error, a term we define below, this is not always the case.

We employ a process-oriented point of view when considering error, as supported by Hofer et al. (2000) and Woods & Cook (1999); we consider “the processes that lead up to success and failure (or potential failures) of a *system*” (Woods & Cook, 1999, p. 30). We define an ITSM process to be in error *if it prevents maximizing the outcomes of interest to the organization.* In corraboration, ITSM is ultimately evaluated in terms of business outcomes (Straub & Nance, 1990; Rockart et al., 1996; Garigue & Stefaniu, 2003). To illustrate, overzealous or awkward security measures might hamper business, in which case ITSM would be in error. Or, an organization may choose to truncate security procedures and absorb attacks, if this is less costly than using the full range of security capabilities to repel attacks. In this case, ITSM would not be in error, despite being hobbled. This separation of the evaluation of ITSM from the practice of ITSM means that an increase in security risk does not necessarily indicate an error. We should point out that our focus is not on malevolent abuse.

3 Methodology

Our research seeks to provide a greater understanding of organizational processes that influence distributed cognition in ITSM. Because ITSM is distributed, and because distributed cognition relies on cues and norms, our primary research question was: *How are cues and norms employed in ITSM?* Our secondary research question was: *How do ITSM challenges foster security risks and/or errors?*

To address these questions, we relied on data from the HOT Admin interviews, which were ongoing at the time of this study (see section 2.1). These *in situ* semi-structured interviews were with a wide variety of IT professionals who devoted at least some of their time to security-related tasks. We should point out that the research questions for the HOT Admin interviews were more general than for this study, in that the interviewer inquired about a wide range of ITSM aspects, from minute routine details to long-term goals. The participants were asked about the nature of security, the challenges they faced, tools used and corresponding likes/dislikes, organizational influences, and security culture, to name a few. Not all topics were discussed at the same level of detail with all of the participants. As is common with semi-structured interviews, the format and number of questions changed according to the particular context. Each interview was conducted by a team of two researchers, in order to reduce interviewer bias, ensure coverage of questions, and allow for probing of details from different perspectives. An interview lasted approximately one hour; each was audio recorded and subsequently transcribed and sanitized to preserve anonymity.

Table 2 Industry sector breakdown

Sector	**Organizations**	**Participants**
Academic	3	18
Consulting	3	3
Financial	2	2
Scientific services	1	2
Manufacturing	1	2
Insurance	1	2
Retail/wholesale	1	1
Technology	1	2
Telecommunications	1	1
Government agency	1	1
Not-for-profit organization	1	1
Total	16	35

The analysis reported here is based on data from 35 IT professionals interviewed for the HOT Admin project. As shown in Table 2, the participants came from 16 unique organizations from 11 different sectors; their positions ranged from IT managers to general IT staff to security administrators and security analysts.

We analyzed the interview data using qualitative description (Sandelowski, 2000). The nature of the data was such that participants could be compared for commonalities only, and frequency analysis (e.g., how often participants performed a given task) was not applicable, because as mentioned above, not all participants were asked the same questions. Instances in the interviews of the concepts in our framework (i.e., cues and norms) were collected and then organized according to various themes.

Our approach to refining the concepts of “cues” and “norms” with respect to ITSM was by means of a modified grounded theory, where, through exercising “theoretical sensitivity,” we accepted the initial concepts of “cues” and “norms”, refining these as needed; the refinements reached “saturation” with the first 26 interviews, while the remaining nine interviews confirmed both the initial categories and their refinements, but did not add to them.

4 Manifestation of Cues and Norms in ITSM

That cues and norms are employed in ITSM as an aspect of distributed cognition is an assumption of this work. Thus, since we adopted the distributed cognition framework, it is not necessarily surprising that we found many instances of cues and norms. However, what is of interest is the particular cues and norms that are employed in ITSM—their widespread use is confirmed by the fact that every interview included at least one example. Thus, the power of the support for the concepts stems from the diversity of the participants and their experiences. We now present a subset of the quotes from select participants to illustrate our findings.

4.1 Cues

Recall that a cue is an occurrence of a signal or clue that participants use to determine when to act and how to act. Our analysis produced ample examples of cues, corroborating Busby's (2001) framework. However, within this category, our data exhibited a distinction that we consider to be important; that is, cues were either *not explicity directed* or *notifications.* To illustrate, one morning P12 checked his email from home to discover that an external organization had intentionally relayed a *notification* to him that suspicious behavior was originating from his network. In contrast, during P9's routine inspection of suspicious behavior, he would check for repeated and rapid attempts to gain access to successive computer network host ports, because where these attempts stop indicates a possible break-in. In the second case, the pattern is not left intentionally, but instead is a kind of footprint or clue of the perpetrator's progress that cues the security practitioner to action.

Gutwin & Greenberg (2000) similarly distinguish between *explicit* and *consequential* communication in their conceptual framework for shared-workspace groupware. Concerning consequential communication, they explain that people pick up information that is unintentionally produced by others as they go about their work. We use the term "not explicitly directed" in order to include situations where a cue is intentionally produced, but not explicitly directed to anyone in particular; this is similar to leaving one's office door open to signal one's availability to anyone present.

4.1.1 Cues that are not explicitly directed

We group *cues that are not explicitly directed* into three (not necessarily exhaustive) categories: (a) *quick view*, (b) *proof of reliability*, and (c) *reminders and hints*. These categories arose from our data.

Quick view: Many security-related cues correspond to departures from normal system behavior (P2, P3, P4, P5, P9, P12, P13, P14, P20, P22, P24, P26) and are provided by tools that monitor IT systems. Interpreting these cues requires knowledge of system behavior (i.e., what behavior is normal versus suspicious). To aid in this process, some tools provide a "*quick view*" (P3, P9, P12, P13, P20, P26). A quick view is a high level indication of the system state that can be read at a glance. The most ubiquitous quick-view tool was email (all participants). As P9, a security analyst in a large academic institution said,

> *I have a variety of email filters that put certain kinds of messages in certain folders, and I'll go look at those folders occasionally, and I'll look at the subject line, and I'll see how many of them there are. So if I'm accustomed to seeing 4 or 5 messages from our wireless management system in the course of the day, and I look in that folder and I see 40, or I see 500, or I see more than I can count, then I know that there is a problem*

Another example of a quick-view tool is the Active Ports program, which monitors all open TCP/IP and UDP ports on a local Windows computer (used by P26). Configuration settings also provide a quick view. P14 described how in his company, it was important to not overwrite existing settings in the Novell Automatic Client Upgrade, because these provide a snapshot of the state of affairs, as well as a cue about the organization's readiness to patch or support particular services.

As the email example illustrates, a quick view may provide a cue to look deeper, if something suspicious appears in the view. For instance, cues to investigate deeper for a security breech include: "*system slowdown, something can't connect; there is an alert of a malfunction; an application is blocked*" (P26); and, on a grand scale, *"disappearance of a critical database"* (P17, P18).

Proof of reliability: The reputation of a tool in itself provides a cue of whether the data produced by the tool is reliable. Some tools have a positive reputation. For instance, the use of a Single Sign On (SSO) system is often taken as a proof of authenticity of the email messages that are delivered through it (P21). Another strategy that participants used for increasing the reliability of tool output is correlating the output with other data sources (P3, P9, P13).

Some tools have a negative reputation, for instance because the tool increases the risk of overlooking critical information by *inserting noisy information into files*, analogous to a word processing program producing an HTML page of "hello world" whose source code is hundreds of lines long (P8); or *loss of data*, such as dropping packets of data when overloaded (P2), or due to the sheer magnitude of data "clogging" a system to the extent that it has to be reinstalled (P24). Other examples of cues that a tool is unreliable include *uncertainty about the effects of a tool*, like a Java client for writing configuration files that don't always take effect (P8, P10); and *misleading messages*, like a diagnostic tool that gives messages that do not correlate with messages from the devices being tested, or with other information (P26).

Reminders and hints: Security practitioners create their own tools, referred to as *scripts*, in response to arising needs that can not be met with existing tools (P2, P3, P4, P9, P12, P13, P20, P22). For example, practitioners mentioned creating scripts for detecting security breaches by recognizing anomalous events, such as an excessive number of emails originating from a single network address (P3). Practitioners also created scripts to collect all the logs from a machine that is suspected of being compromised (P3, P9, P12, P13, P20, and P22).

The naming conventions for these scripts may provide reminders of what they are for, and hints of the kind of information that is needed to run them; that is, cue security practitioners on when and how to act. In particular, a script's name, the names of its input parameters, its error messages, comments inside the script, and when it runs on a service (like the Linux operating system's *cron* utility, which enables scheduling periodic tasks) all provide cues (P12). Scripts are widely used by practitioners; for instance, P12 wrote approximately 2,000 scripts over a 20 year period. P12 had been an IT systems administrator for a large research facility for more than two decades, and had evolved into being the IT security manager. His scripts, which he sometimes submitted to online discussion forums about IT security, informally documented his experiences and acted as reminders for himself and other stakeholders about security issues and what to do about them.

4.1.2 Cues as Notifications

Another form of a cue is a notification, which is a kind of message. Unlike cues that are not explicitly directed, not only does the meaning of a notification have a *constructed sender*, a *constructed receiver*, and other constituents (Fouquier, 1988), the time frame of *now* is an important aspect of the message context. To illustrate, the notification that P12 received from an external organization not only indicated that the problem was happening *now*, but contained the implication that he had to deal with it immediately to avoid having his organization's access to the internet cut off. Our data analysis revealed a rich variety of categories of notification in ITSM including (a) scripted, (b) public, (c) self, and (d) escalated. Note that these categories are not necessarily exhaustive or mutually exclusive.

Scripted: A scripted notification is a pre-planned notification that runs periodically or is triggered by an event. One form is automated scripts, which are programs written by a security practitioner in an interpreted language (e.g., Perl, Shell) to collect and email relevant, machine-generated information to that practitioner (P2, P3, P4, P9, P12, P13, P20, P22). A script may be designed to only send notifications when action is needed, such as the need to investigate a possible intrusion. Alternatively, notifications may be sent in order to maintain a practitioners awareness of a system's routine functioning.

Pre-planned notifications are not always automated. A "script" can be an aspect of a practitioner's informal rules, e.g., a security analyst may routinely copy logs of suspicious activity and paste them into emails to be sent to the network administrators responsible for the affected areas (P24). Reports are another example of scripted notification; they are pre-planned and are written either periodically or in response to some trigger, such as a request for a report. Two examples of reports are a list of required patches (P5, P9, P12) and the count of detected computer viruses (P12).

Public: Many IT security practitioners watch information sources like the Sysadmin, Audit, Network, Security (SANS) Institute, which uses web sites and email lists to generate notifications on security issues. Similarly, vendors generate notifications about products, as well as newly discovered vulnerabilities and their patches. Change management processes, such as prescribed by the Information Technology Infrastructure Library (ITIL) work flow standard, notify the authorized individuals of change requests (P2, P14, P25). P25, an IT security analyst in a financial organization, spoke of change management as if it were a kind of bulletin board where interested people could see and comment on proposed changes.

Self: P26, an IT consultant with hundreds of small clients, used his own records as notifications to himself over time. By means of an audio recorder, he kept a running commentary of his work, which he later transcribed into digital documents that he could search. He used the records to justify his invoices and to revise hypotheses with new data.

Escalated: With independent or semi-independent organizations cooperating over ITSM, it is not a given that the various parties will respond to notifications from security specialists, since the recipients may have conflicting responsibilities or other limitations (P2, P14, P22, P24). Consequently, ITSM often involves notifications that escalate from gentle reminders to the actual disabling of the unresponsive party's access to the organizational network.

In addition to lack of response, escalation of a notification may also occur for other reasons. A *help desk* system may escalate its response to a request, depending on the skill level and authority required to take care of the request (P1, P24). In a similar vein, IT consultants who work with small and medium-sized businesses often have a "*go-to*" person within the client organization (P26). If an incident occurs, the affected person will notify the go-to person, who decides when to notify the consultant. The level of expertise of the go-to person varies from client to client, which influences the degree of operational tasks this person undertakes. For instance, if the level of expertise is high, then the consultant may give the go-to person a CD of diagnostic tools to do a first level of analysis.

As these examples illustrate, the process of escalation is itself a cue to practitioners about when and how to act, while the escalation process is typically governed by policy (a kind of norm), discussed below.

4.2 Norms

Busby (2001) defined norms as "rules of some sort that help make the participants' subtasks consistent with each other." Following Busby & Hibberd (2006), we take a broad view and consider a norm to be an "organizational artifact" used as a resource that is subject to an actor's judgment. Analysis of our data revealed a rich variety of norms. As was the case with cues, to accurately portray our findings, we refined the category to include several subcategories.

4.2.1 Notification procedures

Security practitioners have explicit procedures for notifying stakeholders via reports of ITSM issues (P1, P2, P5, P25). Unsurprisingly, notification norms are often side-stepped in order to take "practical action" (Suchman, 1983). For example, P3's organization regularly used "*back channels*," in which certain persons would contact the security officer directly, usually concerning incidents that had to be handled with sensitivity to legal requirements. P25's organization did not consistently generate ITSM reports, because they were not seen as directly supporting business goals.

As we mentioned above, notification norms typically include escalation procedures. Although these procedures are usually explicit, escalation can blur from a documented procedure to one that is guided by mutual understanding (e.g., the example of the "*go-to*" person).

4.2.2 Consistency

In ITSM, consistency with respect to the IT environment is a critical, explicit norm. For instance, to maintain consistency, IT products must be certified to work with existing applications already adopted by an organization. To illustrate, in P14's organization, Oracle (a relational database management system), Banner (an administrative computing system), and COGNOS (business intelligence and performance management software) had not yet been security-certified to work with FireFox (Web browser); this led to a ban on FireFox and subsequent indignation from some end users. The same organization standardized how assets (e.g., machines) were labeled, making the labels into a representation that was shared by physical security, IT security, accounting, and other departments. This allowed for an asset's location and owner to be quickly identified during a security incident (P6).

Templates: A template is a sort of a *pattern*, used to maintain consistency in ITSM. One example is an *image*, which corresponds to all the software that is needed to set up a computer including applications and user data. Typically, security practitioners apply the same image to many computers; a user then extends the image with personal data. ITSM best practices require that the image be approved by a security practitioner (P14); that is, it is important for ITSM that the image be an explicit norm.

Another template in ITSM is a *job role*. In the management of digital identities, a job role includes its identifying name and the entitlements to access resources that the job responsibilities entail. Job roles can be extended. For example, a new employee would receive a generic "employee" role with the entitlements to access e-mail and the internet. The generic role would be then be extended with the role that the employee is contracted to execute. The employee's personal information plus his or her job roles comprise a *profile*. Typically, there are many requests for exceptional access to resources (P21). If these are not handled with strict adherence to both job roles and procedures for extending and retracting profiles, then the rules controlling access privileges can quickly become unmanageable (P16). P21, a security administrator in a large manufacturing organization, managed 240 job roles. Some of these roles were core roles that she extended in order to define further roles. The main core job role included so many rules that it was unwieldy and inconvenient to edit. Further roles were created by writing rules in separate modules, which were then associated with the core role. To carry on her job, she exercised deep tacit knowledge about the organization.

Audits: Security audits are explicit norms that help maintain organizational consistency. For example, P7 described a security audit checklist that his organization employed when creating new applications to ensure they are compliant with audit standards.

Policy: Another example of an explicit norm for maintaining consistency during ITSM is security *policy*, used to specify security practices within an organization. These practices can range from broad policy statements to detailed guidelines (P5). Designing new policy can be a major undertaking; in a large organization, one page of policy can take a year to develop and approve (P2, P5). (See Werlinger et al. (2009) for a discussion about the interactions that security practitioners have with other stakeholders during the development of policy.)

Other Standards: In addition to formal policies, typically organizations will have less formal standards, such as a standard for log file size and how long the files are kept. Such standards help maintain consistency, for instance by controlling the rate at which the logs are archived, as well as how long they are kept. In one organization, the log size and archive duration for a particular department were relatively small (P14), because the mutual understanding was that security would also keep a copy of the logs for as long as was feasible (P3).

4.2.3 Mutual understanding

We use the term *mutual understanding* to identify a class of norms that involve cooperating parties undertaking a process of establishing *common ground* (Clark, 1996) in order to reach a mutual understanding. We now describe norms that rely on participants having a mutual understanding, including *risk assessment*, *security awareness*, and *professional collaboration.*

Risk assessment: This process of assessing vulnerabilities and risks in the IT infrastructure involves building mutual understanding among the various parties involved. P25 explained that risk assessment requires both top-down support from executive management and bottom-up flexibility from the organizational units. This is necessary in order to balance various points of view as stakeholders strive to bring security risk down to an acceptable level. Nevertheless, in risk assessment, unresolved differences may need to be arbitrated. Furthermore, since risk assessment involves a variety of stakeholders, seeing the policies through to implementation requires top management support. One consultant (P27) would not enter a risk assessment process with a client unless the client's executive management appointed an individual to enforce the resulting policy.

Security awareness: The degree to which stakeholders are aware of security processes is an example of tacit, uncodified mutual understanding. One organization attempted to build a deeper awareness of security by distributing its IT security specialists into various organizational units, as opposed to keeping them centralized (P15). The idea was to "*have a band of security going through.*" Another organization used its logo as the (unchangeable) desktop background image on all company computers in an attempt to build awareness that these machines were not for personal use (P14). Some organizations provide security training for laptop users, including a laptop user guide, walk-through sessions, and establishing an initial test group of security-trained laptop users in the organization (P25).

Professional collaboration: There are clusters of *mutual awareness* and trust between IT consultants, including security specialists (P26, P27). Consultants have their specialties, and will cooperate with other consultants who have complementary niches (P26). In a similar vein, departments within an organization will cooperate with other departments to resolve IT security issues (P1, P2, P4, P5, P14, P15, P22, P24). However, it is important to point out that cooperation is not a given, particularly if there is pressure; involving other stakeholders may require additional overhead.

4.2.4 Transactive memory

Transactive memory is a type of mutual understanding where people in a group know who is responsible for what, and is based on the *"idea that individual members can serve as external memory aids to each other"* (Wegner, 1986). In a small organization, knowing who is responsible for what is rarely an issue. P10, a systems administrator in a modest organization, explained that even though he did not have an FAQ on the organization's website, its members had been there a long time and knew who to contact for help. In complex organizations, on the other hand, this is more challenging, as more stakeholders are involved (P18).

Transactive memory allows members of a group to be able to mobilize each other's knowledge and expertise, which may or may not be tacit. To illustrate, P9, a security analyst for a large academic institution, explained how he worked with a support team to obtain and integrate information critical for upholding security in his organization, and also relied on a network operations center to help monitor the networks. "*If something changes color, they will call the person responsible for it*" (P9). For ITSM to be effective, all these practitioners had to know who was responsible for what, and to be able to activate that person's skill and knowledge by calling on them.

We illustrate a mobilization of tacit knowledge via transactive memory with a quote from P13, a network manager in a large research facility:

> *After years of running this place they figured all this out . . . if this computer crashes, this other one can cycle the power to it remotely so that we do not have to shut down the [main service]. So he had already tried that and he was going to go*

> *and take a look, and I decided to go with him because two eyes are better than one. We got down there and ... we suspected the network, and came to the conclusion that it had to be upstream. And of course we were both wrong because the switch was misleading us. We ended up calling another woman ... for yet a third opinion, and realized that we had not power-cycled the switch. That was her idea. So then we had to go back down, and, once we did that, it fixed it.*

An aspect of effective transactive memory is that it can enhance an organization's ability to practice *mindfulness* (i.e., sensitivity to small but important details) (Weick & Sutcliffe, 2001). For example, one day P4, an IT security manager in a large financial organization, attended a meeting located in another section of the company. The human resources manager noticed him pass by, and made a point of stopping him to convey a report about a "spear phishing" incident. In this case, transactive memory enabled the manager to know who tell about this incident. P4 successfully hunted down the offending server and convinced the corresponding service provider to disable the attacker's web site. This particular incident might not have been detected without the practice of transactive memory; regular phishing incidents are usually detected because they are numerous attacks on a site (P4), but there were only three reports of an attack here.

In summary, we found a rich and varied set of cues and norms in ITSM. To accurately describe them, we refined Busby's (2001) classification to include the concepts described above.

5 Breakdowns of Distributed Cognition

Now that we have described how distributed cognition in general, and various cues and norms in particular, manifest themselves in ITSM, let's look at how some of the ITSM challenges we introduced in Section 2 foster breakdowns in the effective use of cues and norms. The following three incidents were told to us by our participants. The first two incidents indicate error – an ITSM process prevented maximizing the outcomes of interest to the organization – while the third incident indicates an erosion of distributed cognition, but not necessarily error. Here, we also identify the ITSM challenges that fostered the breakdowns of distributed cognition (see Figure 1, bottom).

Example 1: A failure of notification results in a breakdown of distributed work. The minimal requirement for performing distributed work is to notify the team whether a goal has been accomplished (Cohen & Levesque, 1991). As we will see shortly, in one organization this requirement was not realized, because of the challenges of *technology and organization complexity*, *goal-oriented behavior*, and *distributed security* leading to a *failure of notification.* This large organization had the ongoing goal of being prepared for security-related change. Achieving this goal was a continual process, which required that the relevant stakeholders communicate with each other as security related situations arose. (Already we can see the challenges of *complexity of organization* and *distributed security* at play.) In this organization, the failure of some stakeholders to acknowledge an email stalled the installation of an important patch.

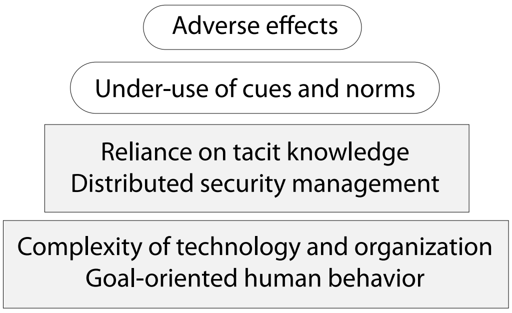
Adverse effects
Under-use of cues and norms
Reliance on tacit knowledge
Distributed security management
Complexity of technology and organization
Goal-oriented human behavior

Fig. 1 Pyramid of ITSM challenges (bottom) in ITSM that culminate in adverse effects (top)

Specifically, P14, a manager for the Microsoft Windows machines, received a phone call from the IT security officer about a serious vulnerability. This meant that patches had to be quickly deployed. Complicating the matter was the fact that the organization's business activity was in a sharp upsurge, and continuity of service was very important. Deployment could break some services, and those people who might be affected needed to be notified so that they could be ready to respond. Consequently, the manager could not perform the deployment until his notification was acknowledged. To identify which services would be affected by the deployment, the manager consulted the Novell documentation ("*I read through everything, I do the Novell administrators' work for them*"). (Now we see *complexity of technology* at play.) He sent an urgent email to the three Novell administrators, but they did not acknowledge it. (The manager considered replying to such an email to be a *norm*, but this norm was not *mutually understood* by the administrators.) The patch deployment was delayed for three working days and a weekend resulting in exposure to the serious vulnerability for five days. The failure to acknowledge the manager's email was a breakdown of *notification* and led to an unnecessary risk for the organization.

We speculate that the Novell administrators were unable keep up with the complexity of the technology (“*I do the Novell administrators’ work for them*”), and that the sharp upsurge in business activity reduced their capacity to handle the IT security aspect of their jobs; that is, they likely experienced conflicting goals. (We see *goal-oriented behavior* at play.) We also speculate that the Novell administrators did not fully appreciate the manager’s expectation of them to reply to his email. The *complexity of the organization* inhibited its stakeholders’ *transactive memory*, which led to a failure of *notification*, resulting in a breakdown of distributed cognition.

Example 2: Lack of norms magnifies a catastrophic breach of privacy. This story, which illustrates the importance of clearly established ITSM norms, was told by P17 and P18 (the organization sector is withheld to ensure confidentiality). Their organization had significant legacy IT systems, as well as many semi-independent projects, IT contractors, and temporary workers. The contractors and temporary workers were responsible for implementing the organization’s security measures, but unfortunately lacked a full appreciation of the necessary measures “*You know, most* techies *are not very good at communicating, so you have to drill it into them.* Some of them prefer to operate on their own” (P17, the business-level IT manager).

One day, a critical database disappeared, which contained confidential information about many clients. It came to light that the database had been mistakenly placed on the outside of the organization’s firewall, and thereby exposed on the Internet. The IT personnel who mishandled the database had violated the organization’s norm to protect the privacy of this database, but the violation had gone unnoticed until the incident. To make matters worse, all the relevant logs had ceased to exist, as described by P18, the IT specialist:

> *Because of the nature of how security was set up in the department, there wasn’t enough log information to identify how the database was attacked and breached, and the data was destroyed out of it. . . . The fact is, it was deleted in the first place because they thought it was an appropriate step to do at some time.*

The security specialists were unable to determine the extent of the damage or when it started. Moreover, recovery was made more difficult because the organization’s IT network supported numerous dependent subnetworks that had to remain operational during the incident recovery.

Both *complexity of technology and organization* contributed to the incompetent handling of the database. The responsible IT personnel in this scenario lacked *mutual understanding* about what to protect and how to do so. Possibly, this IT personnel thought the organization in question was protected by the perimeter of the host organization (*distributed security*). In any case, nothing *cued* the relevant personnel that the current database/firewall setup might be vulnerable. The lack of security expertise in how the logs were set up indicates a failure to mobilize *tacit knowledge* concerning how best to handle the logs. It also illustrates an impoverishment of *transactive memory*: the relevant practitioners were unable to perceive their lack of expertise and ask for the appropriate stakeholder’s assistance. Having learned from their shortcomings, the organization thereafter expended considerable effort to achieve *mutual understanding* about their *norms* by training new staff about privacy and security policies.

Example 3: A workflow standard inhibits both transactive memory and expression of lateral acknowledgement in establishing mutual understanding. Workflow standards (e.g., ITIL) are *norms* that promote an enterprize perspective and *mutual understanding*, whereby individuals are made aware of the effect of their activities on other parts of the enterprize. These change management procedures also provide an opportunity for the enforcement of security reviews or audits within organizations (P2, P14, P25). However, such bureaucratic procedures can also undermine the use of *transactive memory*, which is important for mobilizing tacit knowledge in a distributed scenario – this inhibition of *transactive memory* is a breakdown of distributed cognition. In the words of P14:

> *I can’t just go out and do it now unilaterally, by just talking to [this or that person]. So if I want to make a change on Lotus notes or which affects Lotus notes, I usually pop up to the Lotus notes administrator, see the analyst and say “This is what I want to do. This is the impact. What do you think?” And he, being the guru, he’ll look at that and say yes, excellent deal or no, because. . . Okay, fine. I’ll go away. Now, if he said that’s a great idea, I would say okay, blessed by [the Lotus Notes administrator], thank you, and it would be on the next [PC] image [for configuring new PCs]. Now I can’t do that. Now it goes through all the:* yeah well I have to fill out a change [proposal form] and dah, dah, dah, dah. *And the beauty of it is I don’t even get to speak to the [proposed] change. It’s all in text . . . in a* little *form, which is kind of counter-productive. In my opinion, you should be allowed to present a change [proposal]. Because you only have so*

> *much room to write what you are doing... The only thing is, is that people that don't know what you are doing don't understand, because that's not what their strength is. You know, someone can say: I don't understand, so I'm not putting in a vote for that.*

P14 spoke of days gone by when he would circle his department and talk to key players. "*Now it's all done through committee, so to speak. So, things don't happen necessarily as quickly, and they don't happen necessarily much better.*"

6 Discussion

Cues and norms are an integral aspect of work in distributed environments. We illustrated this in the context of ITSM, which is highly distributed in nature, both with stakeholders across the organization, and with tasks between these stakeholders. As such, it is not surprising we found that cues and norms play a key role in ITSM. Our contribution, however, is that we elaborated the concepts "cues" and "norms" in order to adequately describe our data. In particular, we applied Busby's (2001) understanding of the role of cues and norms in large-scale engineering design to our data, and found that in ITSM cues include (1) cues that are not explicitly directed, such as *quick views*, proofs of reliability, and reminders & hints; and (2) notifications such as scripted notifications, public notifications, notes to self, and escalated notifications. We also found that the use of norms includes (1) notification procedures; (2) methods to maintain consistency, such as templates, audits, policies, and standards; (3) establishment of mutual understanding by means of risk assessment, promotion of security awareness, and professional collaboration; and (4) employment of transactive memory to activate the specialized knowledge and skills of others.

We subsequently illustrated, with three examples, breakdowns in distributed cognition and the corresponding ITSM challenges. Inadequate use of transactive memory (i.e., practitioners failure to adequately know who knows what, and/or activate each other's specialized knowledge) and failure of mutual understanding (i.e., practitioners failure to explicitly acknowledge to each other that they will observe a norm) both seem central in our data. In particular, in all three examples, one stakeholder was unable to realize that his behavior was inadequate to meet ITSM demands, while a second stakeholder was able to realize the inadequacy of the first, but not able to do anything about it. For instance, the Windows security manager was aware of the inadequacy of the Novell administrators' communication practices, the security specialists realized the inadequacy of the ordinary IT technicians' log practices, and the Windows security manager was aware of the inadequacy of a change management process to handle security issues. These findings highlight how the complexity of technology and organization inhibits practitioners from achieving mutual understanding through the practice of explicitly acknowledging to each other that they will observe a norm. In example 1, the Novell administrators never explicitly acknowledged that they would respond to the Windows security manager; in example 2, the IT technicians did not have norms to explicitly acknowledge norms; in example 3, the change management process inhibited lateral acknowledgements.

Altogether, it appears that ITSM challenges not only generate the need to use cues and norms to carry on distributed cognition, but also make the use of cues and norms difficult. Thus, ITSM is susceptible to the under-use of cues and norms; in turn, this erodes the fabric of distributed cognition that ITSM relies on, altogether increasing security risks. This may lead to effects that are adverse to the best interests of the organization (see Figure 1, which summarizes our understanding of the challenges that culminate in adverse effects).

Our work has both theoretical and practical implications. The theoretical implication is to confirm and refine Busby's (2001) framework of cues and norms. In our work, cues and norms gave insight into the process of ITSM, which suggests that the two concepts can be useful for analysis of other, similarly-complex domains. Yet, we found the concepts to be too general to classify further distinctions that are important in ITSM, therefore, to better represent our data, we refined cues and norms as described above. "Cues" and "norms" support the interactions of any distributed network of cooperating actors (including tools), while our subcategories help clarify how they manifest in conditions of technological and organizational complexity, reliance on tacit knowledge, and distributed management.

As with Busby (2001), the practical implications of this work are directed mostly at management to apply heuristics to situations that involve distributed problem solving. These heuristics, for example, could include helping actors to interpret what is required of them by learning which cues are relevant, and what a cue's rationale is; and to help actors be sensitive to both the extent that they use norms and/or to the consequences of a norm's misapplication. In the presence of technological and organizational complexity, reliance on tacit knowledge, and distributed management, upper management's facilitation of transactive memory and ongoing mutual understanding is particularly important.

(See Werlinger et al. (2009) for findings about the development and acceptance of IT security policy.)

Our findings have practical implications for the institutional design aspect of ITSM. Braithwaite (1998) explains that, for institutional design, it is important to consider two kinds of norms that are used to judge trustworthiness: *Exchange* trust norms reflect competence, predictability, consistency, and cautious decision making, while *communal* trust norms emphasize respect for others, sharing of resources, and meeting of others' needs. In the context of exchange trust norms, *trust cues* may provide knowledge about an entity's trustworthiness. For instance, *proof of reliability*, a kind of cue that we identified in our results, is an example of an ITSM trust cue indicating predictability and consistency. We know that ITSM must, more fastidiously than regular IT, keep abreast of advancing technology and evolving organization – what was reliable yesterday is not necessarily reliable today. Therefore, we surmise that, in ITSM, many trust cues must be frequently regenerated. Further complicating the effective use of exchange trust norms is the fact that the various cooperating actors (e.g., humans, tools) are typically specialized, and consequently, the knowledge that they gain about each other to judge trustworthiness is necessarily translated through shared language. This translation, however, may lead to misunderstandings if not carried out properly.

As far as ITSM communal trust norms are concerned, our finding that breakdown of transactive memory was common and central to the three examples regarding breakdown of distributed cognition suggests that social connectedness plays an indispensable role in ITSM. Note that Braithwaite (1998) uses the concept of trust norms in the context of ITSM governance. However, trust in the context of ITSM has several dimensions, including not only governance, but also formal definitions suitable for automated interactions between computers, access control, etc. Therefore, it would be interesting to explore how Braithwaite's characterizations might apply to the other senses of trust.

These practical implications are consistent with the call by Woods & Cook's (1999) for attention to (1) how knowledge that is relevant to the current situation is called to mind, (2) how attentional focus shifts over time, and (3) how balance or trade-offs among multiple interacting goals are accomplished.

7 Limitations and Future Work

The interviews used in this work were primarily obtained from isolated individuals over different organizations; we did not have the opportunity to collect and compare alternate accounts, nor independently verify what was said, let alone perform *in situ* observation. The interviews were not conducted with the sole goal of exploring cues and norms and/or did not always provide a holistic picture. For instance, while it was clear that cues and norms were an integral part of ITSM and their under-use fostered situations of risk, we found less data precisely describing situations of success (possibly because, in ITSM, success is often characterized by lack of a security incident, i.e., typical day to day operations). In general, this work only indicates the existence of cues and norms in ITSM processes, but does not describe the extent and deployment of their manifestation. Also, we did not *observe* instances of distributed cognition. For example, participants told us that they would compare information from different sources in order to confirm their suspicions – a key concept of distributed cognition is bringing disparate information into registration with each other. However, we did not have the opportunity to witness and record such incidents as they unfolded. Nevertheless, our examples (e.g., *proofs of reliability*, *escalation of notification*, *risk assessment*...) may help explain observed instances in future research.

It can be argued that examples can be found for any scheme, and that we found examples of cues and norms in our data is not necessarily surprising. Nevertheless, not only do we consider Busby's (2001) case and approach to be relevant to ITSM, we are encouraged that our examples were rich enough to refine Busby's scheme. Further research is needed to confirm our findings. One interesting avenue corresponds to obtaining *in situ* observational data from multiple stakeholders within an organizations, to obtain a richer picture of distributed cognition in ITSM.

Acknowledgments

We would like to thank the many generous IT security professionals who participated in our research. This work has been supported by the Canadian NSERC Strategic Partnership Program, grant STPGP 322192-05.

References

Ackerman, M. S. & Halverson, C. (2004). Organizational memory as objects, processes, and trajectories: An examination of organizational memory in use. *CSCW*, *13*, 155–189.

Beznosov, K. & Beznosova, O. (2007). On the imbalance of the security problem space and its expected consequences. *Information Management & Computer Security*, *15*(5), 420–431.

Botta, D., Werlinger, R., Gagné, A., Beznosov, K., Iverson, L., Fels, S., & Fisher, B. (2007). Towards understanding IT security professionals and their tools. In *Proceedings of SOUPS*, (pp. 100–111).

Braithwaite, V. (1998). Communal and exchange trust norms: Their value base and relevance to institutional trust. *Trust and governance*, *1*, 46–74.

Busby, J. & Hibberd, R. (2006). The role of coordination of organizational artefacts in distributed cognition, and their failure in maritime operations. *Le Travail Humain*, *69*(1), 25–48.

Busby, J. S. (2001). Error and distributed cognition in design. *Design Studies*, *22*, 233–254.

Chandrasekaran, B. (1990). Design problem solving: A task analysis. *AI Magazine*, *11*(4), 59–71.

Chebrolua, S., Abraham, A., & Thomas, J. (2005). Feature deduction and ensemble design of intrusion detection systems. *Computers and Security*, *24*(4), 295–307.

Clark, H. H. (1996). *Using Language.* Cambridge, England: Cambridge University Press.

Cohen, P. & Levesque, H. (1991). Teamwork. Technical report, SRI, Menlo Park, CA.

Fouquier, E. (1988). Figures of reception: Concepts and rules for a semiotic analysis of mass media reception. *Int. J. of Research in Marketing*, *4*(4), 331–348.

Fuchs, L. & Pernul, G. (2007). Supporting compliant and secure user handling – a structured approach for in-house idm. In *Proceedings of ARES*, (pp. 374–384).

Gagné, A., Muldner, K., & Beznosov, K. (2008). Identifying differences between security and other IT professionals: a qualitative analysis. In *Proceedings of HAISA*, (pp. 69–80).

Garigue, R. & Stefaniu, M. (2003). Information security governance reporting. *EDPACS*, *31*(6), 11–17.

Goel, V. & Pirolli, P. (1992). The structure of design problem spaces. *Cognitive Science*, *16*(3), 395–429.

Goodall, J. R., Lutters, W. G., & Komlodi, A. (2004a). I know my network: Collaboration and expertise in intrusion detection. In *Proceedings of CSCW*, (pp. 342–345).

Goodall, J. R., Lutters, W. G., & Komlodi, A. (2004b). The work of intrusion detection: Rethinking the role of security analysts. In *Proceedings of AMCIS*, (pp. 1421–1427).

Gutwin, C. & Greenberg, S. (2000). The mechanics of collaboration: developing low cost usabilityevaluation methods for shared workspaces. *Proceedings of IEEE Int. Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises*, 98–103.

Haber, E. M. & Bailey, J. (2007). Design Guidelines for System Administration: Tools Developed through Ethnographic Field Studies. In *Proceedings of CHIMIT*, (pp. 1–9).

Hawkey, K., Muldner, K., & Beznosov, K. (2008). Searching for the Right Fit: Balancing IT Security Model Trade-offs. *IEEE Internet Computing*, *12*(3), 22–30.

Hofer, T. P., Kerr, E. A., & Hayward, R. A. (2000). What is an error? *Effective Clinical Practise*, *3*(6), 261–269.

Hutchins, E. (1995). *Cognition in the Wild.* Cambridge, MA: MIT Press.

Kandogan, E. & Haber, E. M. (2005). Security administration tools and practices. In L. F. Cranor & S. Garfinkel (Eds.), *Security and Usability: Designing Secure Systems that People Can Use* (pp. 357–378). O'Reilly Media, Inc.

Kraemer, S. & Carayon, P. (2007). Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists. *Applied Ergonomics*, *38*, 143–154.

Maglio, P. P., Kandogan, E., & Haber, E. (2003). Distributed cognition and joint activity in collaborative problem solving. In *Proceedings Conference of the Cognitive Science Society.*

Polanyi, M. (1966). *The Tacit Dimension.* Garden City, New York: Doubleday & Company, Inc.

Poole, M. S., Seibold, D. R., & McPhee, R. D. (1985). Group decision-making as a structurational process. *Quarterly Journal of Speech*, *71*, 74–102.

Rockart, J., Earl, M., & Ross, J. (1996). Eight imperatives for the new IT organization. *Sloan Management Review*, *38*(1), 43–55.

Salomon, G. (1993). *Distributed Cognitions: Psychological and Educational Considerations.* Cambridge University Press.

Sandelowski, M. (2000). Whatever happened to qualitative description? *Research in Nursing & Health*, *23*(4), 334–340.

Schultz, E. E. (2007). Computer forensics challenges in responding to incidents in real life setting. *Computer Fraud & Security*, *12*, 12–16.

Siegel, D. A., Reid, B., & Dray, S. M. (2006). IT Security: Protecting Organizations In Spite of Themselves. *Interactions*, 20–27.

Simon, H. A. (1973). The structure of ill structured problems. *Artificial Intelligence*, *4*(3), 181–201.

Straub, D. & Nance, W. (1990). Discovering and disciplining computer abuse in organizations: a field study. *MIS Quarterly*, *14*(1), 45–60.

Suchman, L. (1983). Office procedure as practical action: models of work and system design. *Transactions on Information Systems*, *4*(1), 320–328.

Wegner, D. M. (1986). *Transactive memory: A contemporary analysis of the group mind.* In B. Mullen and

G. R. Goethals, Eds., Theories of Group Behavior.

Weick, K. & Sutcliffe, K. (2001). *Managing the unexpected: assuring high performance in an age of complexity.* Jossey-Bass.

Werlinger, R., Hawkey, K., & Beznosov, K. (2009). An integrated view of human, organizational, and technological challenges of IT security management. *J. of Information Management & Computer Security, 17(1)*, 4–19.

Werlinger, R., Hawkey, K., Botta, D., & Beznosov, K. (2009). Security practitioners in context: Their activities and interactions with other stakeholders within organizations. *Int. J. of Human-Computer Studies*, 1–41.

Woods, D. & Cook, R. (1999). Perspectives on Human Error: Hindsight Biases and Local Rationality. *Handbook of Applied Cognition*, 141–71.

Zhang, J. (1998). A distributed representation approach to group problem-solving. *J. of American Society of Information Science*, *49*(9), 801–809.