Value Tensions in Design: The Value Sensitive Design, Development, and Appropriation of a Corporation's Groupware System

Jessica Miller
Computer Science and Eng.
University of Washington
Seattle, WA, USA
jessica@cs.washington.edu

Batya Friedman
The Information School
University of Washington
Seattle, WA, USA
batya@u.washingon.edu

Gavin Jancke
Microsoft Research
Redmond, WA, USA
gavinj@microsoft.com

Brian Gill
Mathematics Department
Seattle Pacific University
Seattle, WA, USA
bgill@spu.edu

ABSTRACT

We report on the value sensitive design, development, and appropriation of a groupware system to support software engineering knowledge sharing. Usage data (5,965 visitors) and semi-structured interviews (18 individuals) suggest the methods employed were successful in addressing value tensions, particularly with respect to privacy, awareness, and reputation. Key contributions include: (1) "proof-of-concept" that established Value Sensitive Design principles and methods can be used to good effect for the design of groupware in an industry setting, (2) a new design method for addressing value tensions, Value Dams and Flows, and (3) demonstration of the co-evolution of technology and organizational policy.

Categories and Subject Descriptors

H.5.3 [**Group and organizational interfaces**]: Computer supported cooperative work, Evaluation/methodology

General Terms

Design, Human Factors

Keywords

design method, Value Sensitive Design, value dams and flows, human values, groupware, organizational culture, emergent issues, privacy, awareness, reputation

1. INTRODUCTION

Over the past two decades, designers, researchers, and engineers have built numerous information systems to support collaboration and cooperation in and amongst groups. Though the potential benefits of such systems are quite attractive, many of these technologies have met with considerable social and organizational challenges [11, 14, 15] – frequently with human values at their crux.

As well, there has been increasing attention in the computer-supported cooperative work and the broader human-computer interaction communities to the place of stakeholders' social, moral, and political values in system design and development. Implicit in this work, the definition of *value* goes beyond that of economic worth to embrace what a person or group of people consider important in life. One strand of research focuses on understanding the meaning and components of specific values (e.g., informed consent, fairness, security, trust, privacy) with respect to technology [5, 6, 13, 16]. A second strand investigates stakeholders' views about and experience of values with respect to technology [8, 14, 15, 18, 19]. A third strand proposes technological designs to support specific human values [1, 2]. Yet a fourth strand investigates methods to design for values [3, 4, 9, 10].

That said few practical methods exist to address tensions among a diverse group of values as they unfold during the design and deployment process. Moreover, with the exception of Participatory Design, the vast majority of value-oriented research that does exist has been carried out in laboratory or short-term field study settings. While providing valuable results, such work needs to be complemented by longer-term design research in industry settings that engages the complexity of human values in situ. The work reported here begins to address both concerns.

Specifically, we set forth to answer the following questions: What existing principles and methods for discovering and supporting values can be carried over from research settings to industry's faster-paced, bottom-line orientation? What new methods can be envisioned to help designers navigate value tensions among system features in a principled way? How might design methods be leveraged not only to develop technical features but also appropriate organizational policies to support the system and its users? Over time, what unexpected stakeholders and values-related issues emerge and how might designers account for these?

To gain traction on such challenging questions, we believed we needed to (1) design, develop, and deploy a system from scratch; and (2) employ a design approach that accounts for human values. We identified the following additional requirements for the system: defined to meet a perceived organizational need; documented values implications; and confined to an organizational setting and user population that could reasonably be studied. Thus, we chose to conduct our design and development work with the CodeCOOP, a groupware system targeted for several hundred engineers in the research arm of large software engineering organization (which we refer to as "LEO-R"). In terms of design approach, we wanted a robust methodology to help structure our work. Here we utilized Value

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GROUP'07, November 4–7, 2007, Sanibel Island, Florida, USA.

Copyright 2007 ACM 978-1-59593-845-9/07/0011...$5.00.

Table 1. LEO-R's and VSD's Collaborative Design Process of the CodeCOOP. Though displayed as a list, design events often occurred concurrently. For LEO-R Processes, L indicates their usual design process; for VSD Processes, C indicates conceptual investigations, T technical investigations, and E empirical investigations.

Phase	Major Design Events	Conducted By	
		LEO-R	VSD
Pre-Collaboration	Decision: Build a knowledge sharing tool for LEO-R – the CodeCOOP	L	
Designing and Developing the CodeCOOP (16 months)	Conduct stakeholder and harms/benefits analyses		C
	Solicit informal feedback and design critiques from LEO-R groupware researchers	L	C
	Paper prototyping	L	
	Engineer software prototype of CodeCOOP	L	T
	Survey: Pre-development work practices, views, and values (LEO-R)		E
	Value Dams and Flows analyses		E
	Engineer operational version of CodeCOOP	L	T
Organizational Iteration and Deployment (8 months)	Decision: Include and index company-wide engineering-related DLs	L	
	Decision: Design CodeCOOP repository to be IP-safe	L	
	Functionality and interface testing	L	
	Deployment I: Engineers and researchers in research arm (LEO-R)	L	
	Initiate LEO-R code contribution contest with monetary reward	L	E
	Survey: Formative assessment (LEO-R)		E
	Deployment II: Engineers corporation-wide (LEO)	L	
Assessment and Design Reflection (3 months)	Usage data analysis	L	E
	Interview: Current state assessment (LEO)		E
	Design reflection		C, E

Sensitive Design, an established design approach that accounts for human values in a principled and comprehensive manner throughout the design process. For purposes of this work, our overarching design challenge was to proactively address value tensions in both system features and organizational policies such that value tensions did not undermine system appropriation.

In this paper, we first provide a brief background on the CodeCOOP and our partnership with LEO-R, value-oriented groupware challenges, and Value Sensitive Design. Next we discuss the design and development of the CodeCOOP including an innovative method we developed to address tensions among competing values and organizational policies in the design process. Then we report on the CodeCOOP's deployment to the original target population, LEO-R employees engaged in software engineering, as well as to an expanded population, all software engineering employees in LEO-R's parent corporation. We assess and discuss the CodeCOOP in terms of system appropriation. Finally, we conclude with a statement of our contributions and provide some future directions.

2. THE CODECOOP AND OUR PARTNER ORGANIZATION

The CodeCOOP was originally intended for employees engaged in software engineering at LEO-R, a computing research organization consisting of several research labs worldwide. LEO-R is organized by research groups typically composed of 6-12 researchers and 1-2 engineers. Acknowledging the physical and logical separation among employees engaged in engineering, LEO-R management decided to design and build a system – the CodeCOOP – to support two key goals: (1) to increase and make more efficient information (i.e. advice, code) sharing among engineers and researchers who engineer, and (2) to support community among engineering employees. Key features of the resulting internal web application would eventually include discussion forums, a code repository, a searchable archive of discussion lists, and full-text search over all content.

At the onset, LEO-R management was aware that many such systems (both at LEO-R and in other organizations) had been built to support similar constructive goals and yet failed due to a number of social and organizational challenges. Thus from the start, the Value Sensitive Design team was invited to participate as co-designers to proactively address values issues around the CodeCOOP.

The CodeCOOP's design, deployment, and assessment process, spanning 2 years and 3 months, integrated Value Sensitive Design processes with LEO-R design and development processes. The major events of our collaborative design process are detailed in Table 1. Typically design events overlapped, spanning several weeks or months depending on the event.

3. VALUE-ORIENTED GROUPWARE CHALLENGES

Groupware systems have met, at times, with considerable social and organizational challenges that at best stunt system appropriation and at worst undermine appropriation in its entirety. Value tensions often lie at the core of these challenges and may surface in a number of ways [11, 14, 15]. In some cases, values held by the group conflict with those held by the individual. For example, open calendaring systems have supported group awareness of others' activities and presence at the expense of individuals' privacy [15]. In other cases, large inequities exist among those who contribute to the system and those who benefit from the system. A case in point, knowledge-sharing systems (such as that of the CodeCOOP) are dependent on the good will of a critical mass of key "knowledgeable" participants to contribute content, advice, and time, yet typically these individuals do not receive benefit from or recognition for their participation in the system [11]. In still other cases, groupware systems support values that are at odds with those fostered by the organization's culture and reward structure. Such conflict is exemplified by Orlikowski's study of the failed adoption of Lotus Notes in a consulting firm whose organizational structure rewarded competition rather than collaboration [14].

When value tensions go unaddressed, consequences can range from lack of appropriation by disadvantaged groups to more severe consequences such as system sabotage. As a case in point, the Virtual Kitchen [12], built to increase sociality among employees by linking several kitchens at the workplace with continuous video and audio, was sabotaged by stakeholders (presumably with privacy concerns) who placed notes in front of cameras and, at times, completely disconnected the system.

Participatory Design represents one important value-oriented design methodology that has been applied to the design of groupware systems [9, 10]. Developed in Norway in the 1970s as a result of labor laws entitling workers along with management to co-determine technologies to be introduced into the workplace, Participatory Design has strong commitments to workplace democratization and human welfare. Correspondingly, traditional Participatory Design emphasizes substantive and sustained user participation. In the case of the CodeCOOP LEO-R's management was unable to commit to the level of participation we believed necessary to practice Participatory Design. As such, we determined that Participatory Design was not a good fit for this project; others have reported similar obstacles for utilizing Participatory Design in industry settings (e.g., [10]).

4. VALUE SENSITIVE DESIGN

We ground our work in Value Sensitive Design [4]. To date, Value Sensitive Design has been applied to a range of technologies, including human-robot interaction, large displays, open source software systems, network browser security, and urban simulation [4, 5, 7], but at the time we began had not been applied to groupware or in an industry context.

Value Sensitive Design provides an interactional approach to designing for values in technology. This approach complements previous research reporting complex social and social-technical interactions around groupware. Methodologically, Value Sensitive Design offers a three-part framework in which conceptual, empirical, and technical investigations are applied iteratively and integratively, with new investigations building on results from earlier ones. *Conceptual investigations* focus on values discovery and informed analyses of these values and potential value tensions. *Empirical investigations* serve to refine and expand key values identified in the conceptual investigations and to assess stakeholders' experience of the value-oriented features of a system at various stages of the design, development, and deployment process. *Technical investigations* focus on the design of the technology itself. In the CodeCOOP work reported here, we used all three investigations, iteratively and integratively. Moreover, we integrated these investigations with our partner organization's on-going design processes.

A second key feature of Value Sensitive Design entails identifying direct and indirect stakeholders. Direct stakeholders (i.e. users) interact directly with the system. Indirect stakeholders do not directly interact with the system but are affected by the system or its output. Stakeholders are identified by roles they have with respect to the system or organization. Job title, willingness to contribute to the system, and type of system use are likely criteria for identifying stakeholders of collaborative technologies. During our conceptual investigation, we identified CodeCOOP stakeholders with respect to their interaction with the system as well as their job title at LEO-R.

Another implication of Value Sensitive Design's interactional theory is the working proposition that the co-evolution of technology and policy may be advantageous for addressing value issues in design. Prior work on integrating privacy commitments into an open-source license established this linkage between policy and technology development [7]. With the CodeCOOP, we explicitly sought to co-evolve technology and organizational policy and used the same analysis method for both.

Finally, prior work on bias in system design [6] alerted us to the possibility of emergent value issues as the system was appropriated over time and changes occurred in population and organizational culture. Thus, with the CodeCOOP we paid attention to changes in the deployment population, organizational practices and norms, and technical functionality.

5. DESIGN AND DEVELOPMENT: UTILIZING ESTABLISHED VALUE SENSITIVE DESIGN METHODS

In designing and developing the CodeCOOP, we had all the usual groupware goals of supporting work practice, community, sharing knowledge, usability, and functionality. To that, and we believe largely unique to this work, we attended to value tensions in the design. Here, we report on the novel aspects of our approach that addressed potential value tensions.

We began our design of the CodeCOOP by applying established Value Sensitive Design methods. Specifically, we conducted conceptual investigations to identify stakeholders and understand key values. Then, to refine and validate those results we conducted an empirical investigation of stakeholder views and values by means of an anonymous, online survey.

5.1 Conceptual Investigations

During our conceptual investigations we applied two Value Sensitive Design methods and iterated on the results. Informal feedback from five LEO-R expert groupware researchers with experience designing and deploying systems in the organization provided helpful information with respect to LEO-R culture.

The first method entailed stakeholder analysis. We identified CodeCOOP direct stakeholders by their interaction with the system. Four interaction roles were identified: searcher (someone who searches the system for information), querier (someone who submits questions to the system via a discussion thread), answerer (someone who answers discussion thread questions), and proactive poster (someone who posts code or information without being prompted). As typical with this type of role-based analysis, a single user may engage in multiple interaction roles. In terms of indirect stakeholders, we identified LEO-R managers and executives. Given the need to scope our work to a reasonable size, our initial analyses of indirect stakeholders did not extend beyond LEO-R employees.

Building on our stakeholder analysis, we then employed a second Value Sensitive Design method: identification of potential benefits and harms and corresponding values. In order to conduct the harms/benefits analysis, we reviewed relevant literature about collaborative technologies developed and deployed for similar purposes and in similar contexts, examined other existing groupware systems both within and outside of the organization, and discussed our analyses with the LEO-R expert groupware researchers. Specifically, for each stakeholder type established during our stakeholder analysis, we sought to identify potential benefits and harms that could result from system use or organizational policies to support the system For example, a CodeCOOP user in the role of a querier could benefit from the system if asking a question leads to obtaining an answer, or could be harmed if the question is perceived by the community to be too obvious. After enumerating benefits and harms, we mapped potential underlying values to each. In the previous example potential values include collaboration and reputation, respectively. Other potential values implicated by the CodeCOOP included privacy, awareness, anonymity, trust, attribution, and time.

5.2 Stakeholder Views and Values

In order to refine and validate the results of our conceptual investigations, we then solicited views and values of direct stakeholders through an anonymous, online survey. Specifically, we sought stakeholder perspectives on potential technical features, managerial policies, and corresponding value implications that were generated during the conceptual investigations. A recruitment email was sent to all LEO-R employees in engineering-related positions at English speaking campuses (322 total), and 72 responses were received (22% response rate, 92% male, 7% female). The survey addressed current work practices, value-oriented issues, as well as potential organizational policy to support contributions to the CodeCOOP. The survey instrument included both structured questions (asking participants to rate their agreement with statements on a 5-point "strongly disagree" to "strongly agree" Likert scale) as well as open response questions that gave participants the opportunity to provide unstructured feedback. Table 2 provides a summary of selected results for the structured value-related questions, grouped by whether the question addresses a potential harm or a potential benefit. Additional value-related questions regarding issues such as time, collaboration and building community were also asked, but are not reported due to space limitations. Table 3 provides a summary of selected results for the structured questions on managerial policies to encourage contribution to the CodeCOOP.

We also developed interview instruments for and intended to collect data with indirect stakeholders of the system, in particular, LEO-R managers and executives. However, due to indirect stakeholders' time constraints, we were not able to do so. How to effectively incorporate views and values of indirect stakeholders who work at a high-level within the organization remains an important open methodological question.

6. DESIGN AND DEVELOPMENT: EXTENDING VALUE SENSITIVE DESIGN WITH VALUE DAMS AND FLOWS

We then turned to the problem of translating the results of our direct stakeholder survey into specific, implementable design features. Furthermore, we hoped to do so in a way that was systematic and empirically-informed as well as compelling and tangible. Since we knew of no such existing method, we developed a new design method, Value Dams and Flows.

6.1 The Value Dams and Flows Method

Stated most generally, the Value Dams and Flows method works by (a) avoiding features that even a small number of stakeholders view as particularly problematic, (b) identifying and designing for values stakeholders wish to see the system embody, and (c) systematically addressing values-oriented design tradeoffs. The details follow.

Value Dams refer to technical features or organizational policies that are strongly opposed by even a small set of stakeholders. As such, value dams in extreme cases can undermine system appropriation. Recall the opposition of a small but vocal group of stakeholders to video cameras in the Virtual Kitchen project [12] due to privacy concerns. In addition to the pragmatic rationale for attending to value dams, there is also an ethical perspective – it is good practice to consider the rights and harms of persons in the minority.

That said, a designer should also understand the needs and desires of the majority. *Value Flows* refer to technical features or organizational policies that, for value reasons, a large percentage of stakeholders would like to see included in the overall system, even if the features or policies are not absolutely necessary for successful appropriation. Consideration of value flows can help the designer discover features or policies to draw stakeholders to the system or to make systems with uneven benefits more attractive to those who contribute.

Once dams and flows have been identified, designers must find a way to balance addressing value dams which are of concern to even a small set of stakeholders with value flows supported by the majority of stakeholders. We recognize that these two ideas may frequently be at odds (e.g., privacy and awareness), so we use the CodeCOOP as an example to explore how they can be used together to address values-oriented design tradeoffs.

6.2 Applying the Method to System Features

As a first step, we used the results of our online survey to identify which of the harms and benefits were perceived by stakeholders as value dams and flows. We examined the survey questions concerning potential *benefits* and ranked them by the percentage of participants who agreed or strongly agreed. A natural break in the percentages occurred around 50%. Thus we used the 50% or above as a threshold for establishing a value flow; benefits which exceeded this threshold are flagged with a † in Table 2. This threshold also has the intuitive appeal of identifying any potential benefit the majority of respondents agreed with as a value flow.

Analogously, for value dams, we examined the questions concerning potential *harms*, and ranked them according to the

Table 2. Pre-development Selected Survey Results: Value-oriented Technical Features. N=72 respondents, self categorized as expected contributors (C, N=54) or non-contributors (NC, N=17). An * indicates harms ≥ 10% value dam threshold, and a † indicates benefits ≥ 50% value flow threshold.

Value	Harms-related Questions	Percent Strongly Agree		
		C (N=54)	NC (N=17)	Total (N=72)
Privacy	I would feel like my privacy is being compromised if the system logged....			
	...what and how I searched.	9%	29% *	14% *
	...how I query the system.	7	24 *	11 *
	...information about my answers.	4	24 *	8
Reputation	Contributing incorrect or buggy code could harm my reputation.	9	24 *	13 *
	Posting an obvious or simple question could harm my reputation.	6	12 *	7
	Answering a question poorly could harm my reputation.	2	24 *	7
	Contributing to the CodeCOOP could cause people to criticize my postings.	4	12 *	6
	I would be offended if my answers or posts got a poor public ranking.	0	12 *	3
Trust / Anonymity	I would be less inclined to...			
	...trust an anonymous reply to a question.	17 *	13 *	15 *
	...trust an anonymous post.	15 *	12 *	14 *
	...answer an anonymous question.	9	12 *	10 *

Value	Benefits-related Questions	Percent Agree or Strongly Agree		
		C (N=54)	NC (N=17)	Total (N=72)
Awareness	I would like to...			
	...be able to know how frequently my code and ideas are used.	65†	83†	76†
	...see how my peers have ranked posts as I am searching the CodeCOOP.	72†	41	64†
	...know who most frequently posts and answers questions.	56†	29	49
	...know who has looked at my posts and answers.	43	35	40
Reputation	Answering a question well could benefit my reputation at LEO-R.	83†	53†	75†
	Contributing good code will benefit my reputation at LEO-R.	78†	53†	71†
Anonymity	I would like to be able to post			
	...questions anonymously to the CodeCOOP.	30	41	32
	...proactively to the CodeCOOP anonymously.	20	29	22
	...answers to the questions anonymously to the CodeCOOP.	24	24	22

percentage of participants who strongly agreed. However, no natural break was identified (i.e. percentages for harms-related questions ranged evenly between 0% and 17%). Thus, we drew

on the Kitchen Project research [12], the median value of 9.7%, and our own intuitions to set a threshold of 10% or above for a potential harm to be identified as a value dam. While a threshold of 10% in some respects is arbitrary, we believe it worked well in this project. Value dams identified based on this threshold are flagged with an * in Table 2.

In addition to this overall dam and flow analysis (Table 2, "Total" column), we also conducted separate analyses for self-categorized likely contributors and non-contributors to the system (Table 2, "C" and "NC" columns, respectively). Previous groupware literature indicates a need to adequately support those who do the work [11] so as we made decisions about technical features, we gave priority to contributors' value dams and flows. At the same time, we considered ways to address dams of non-contributors in an effort to alleviate their concerns and draw them to the system.

In order to make decisions about specific technical features, we considered the value dams and flows together. For example, one common value tension faced by designers entails balancing needs for privacy with awareness. In the context of the CodeCOOP, when considering potential privacy-related harms, logging both searches and queries arose as value dams in our survey with 11% or more of respondents strongly agreeing that each of these compromised their privacy. When considering awareness-related potential benefits of the system, a solid majority of participants agreed or strongly agreed that they would like the system to report *how often* their contributions are used (76%) and how their peers ranked their posts (65%), making these two features value flows. However, respondents were not as enthusiastic about knowing *who* used their posts and answers (40%). Thus, to mitigate the privacy-related value dams while still taking advantage of the awareness-related value flows, we determined not to log or report *who* searches or queries, but to log and report *frequency* of code use and implement content ranking.

Another value tension that arose in the design of the CodeCOOP entailed balancing the potential reputation-related benefits and harms that one could incur by posting to the system. Again, considering value dams and flows together gave us insight into how to address these tensions. Several reputation-related harms were identified as value dams, particularly with respect to non-contributors. Originally we thought we could build in a mechanism to allow users to limit the risk to their reputation by allowing anonymous postings. However, most harms associated

Table 3. Pre-development Selected Survey Results: Managerial Policies about the CodeCOOP. N=72 respondents, self categorized as expected contributors (C, N=54) or non-contributors (NC, N=17). An * indicates a value dam (i.e. ≥ 10% strong disagreement), and a † indicates a value flow (i.e. ≥ 50% agreement).

Regarding policies to encourage people to contribute to the system, I would be comfortable if management…	Percent Strongly Disagree			Percent Agree or Strongly Agree		
	C (N=54)	NC (N=17)	Total (N=72)	C (N=54)	NC (N=17)	Total (N=72)
...considered contributions to the knowledgebase in employees' annual evaluations.	6%	12% *	9%	54% †	41%	50% †
...management recognized top contributors with physical rewards (money, movie tickets, etc).	8	6	9	52†	41	49
...management recognized top contributors to the knowledgebase by way of public announcements.	9	6	10 *	23	29	26
...management had no role in encouraging use of the system.	11 *	0	9	35	47	39

with anonymous posts were value dams both overall and to contributors (whose harms and benefits we gave priority), with 10% or more of respondents strongly agreeing that they would be less inclined to trust anonymous replies (15% Total; 17% C) and anonymous code posts (14% Total; 15% C) as well as less inclined to answer anonymous questions (15% Total; 9% C). Furthermore, there was not particularly strong support for being able to post anonymously, with none of the three anonymity-related benefits qualifying as a value flow (see Table 2). Thus, though the survey indicated concern about reputation-related harms, it also indicated that anonymous posting was a poor mechanism to address such concerns. Hence, we did not implement anonymous posting in the CodeCOOP. Instead, since many of the reputation-related harms concerned contributions of poor quality, we designed the CodeCOOP to allow contributors to fully edit their posts to correct errors or improve quality should they wish to do so.

6.3 Applying the Method to Organizational Policy

Recognizing groupware systems do not exist independent of complex social systems, we aimed to recommend to LEO-R management a few appropriate policies to support the CodeCOOP and its contributors. To do so, we applied the same value dams and flows method we developed to address design tradeoffs for technical features.

Using the same thresholds, we turned to the survey data about managerial policies for the CodeCOOP. The survey questions about policies were asked in terms of level of comfort with proposed policies. Therefore, if ≥ 10% of respondents strongly disagreed they would be comfortable with a policy, we considered that policy a value dam (flagged with a * in Table 3); if ≥ 50% agreed or strongly agreed with a policy we considered it a value flow (flagged with a † in Table 3). As we did for technical features, we conducted separate analyses for self-categorized contributors and non-contributors to the system (Table 3, "C" and "NC" columns, respectively).

Again, in order to make decisions about policies to support CodeCOOP contributions, we considered the dams and flows together. The one dam for contributors entailed no managerial involvement in encouraging the use of the system, with 11% of contributors strongly disagreeing. Thus, managers should do *something* to recognize contribution. One proposed approach involved public recognition of CodeCOOP contributors. However, this policy was an overall dam (10% strongly disagreed with this policy), so we ruled out this option. Of the other proposed managerial policies to support CodeCOOP contribution two were seen by contributors to be value flows: (1) recognition in annual evaluations (54% agreeing or strongly agree) and (2) physical (i.e. monetary) rewards (52%). Therefore, the Value Sensitive Design team recommended to LEO-R management two policies for consideration: (1) positively reflect contributions in annual evaluations (since 12% of non-contributors strongly disagreed with this policy, making it a dam, we recommended that management not penalize non-contributors) and (2) offer physical (i.e., monetary) rewards to encourage contribution. LEO-R management decided the system was too new to formally consider recognition of contributions in annual evaluations, and elected to follow the second recommendation, offering the top three code contributors at LEO-R monetary rewards annually.

7. ORGANIZATIONAL ITERATION AND DEPLOYMENT: EMERGENT VALUE ISSUES AND POPULATION

This next phase – of organizational iteration, deployment, and system appropriation – was marked by a shift in control of system design from the Value Sensitive Design team to the LEO-R's CodeCOOP team. Based on prior research [6, 14], we were aware that new stakeholders and unanticipated values issues might emerge as the system was deployed and integrated into the organization. Such was the case with at least three executive decisions made by the LEO-R's CodeCOOP team that we describe briefly here. Note that during this phase we conducted a survey with the original target population; however, we do not report those results here due to space limitations.

The first executive decision entailed populating the CodeCOOP corpus with over 500,000 pre-existing engineering-related distribution list (DL) emails. Distribution lists (an email alias used to contact a group of people) have been in use at the organization for over ten years as a means to share technical information and receive help. Although a publicly accessible DL archive existed, there was no means to search it easily. Thus, incorporating these pre-existing DLs into the CodeCOOP served two purposes: (1) it rapidly populated the CodeCOOP with a large body of meaningful content, and (2) it provided free text search over the DL archive. However, from a values perspective a tension exists among such benefits and the fact that contributors to DLs had never consented to or expected their email contributions to be indexed or accessed in this way. A second potential tension concerns the communities formed around DLs that might not wish to be exposed as visibly to the entire organization. Therefore, from a Value Sensitive Design standpoint, contributors to DLs and the communities formed around them emerged as unexpected indirect stakeholders of the system. Even though we identified the need to solicit stakeholder views on the integration of DLs into the CodeCOOP, organizational constraints prevented us from doing so at this time.

The second executive decision with emergent value implications involved the intellectual property of code in the CodeCOOP's repository and brought a different group of unexpected indirect stakeholders into the mix: LEO's legal department. Specifically, "IP Safety" is a major concern for software engineering corporations due to the abundant supply of "free" code available on the Internet. Should such code become a part of a commercial product, there may be complex legal implications for the corporation. Thus, management made a decision to limit the code contributed to the CodeCOOP's repository to code completely owned by LEO. With this decision comes the corresponding need for a mechanism to ensure that contributed code would be indeed IP-safe. To achieve this, the organization's legal department generated terms for what would constitute IP-safe code (e.g., that the code was "original work written entirely by LEO employees", contained no third party components, etc.) and those terms were integrated into the process for uploading code into the repository. Thus the responsibility to ensure that the code is IP-safe fell to the contributor of the code.

At this juncture, once the IP-safe terms were generated and the software tested both for functionality and usability, the CodeCOOP was deployed. The system was officially announced by email to all employees worldwide at LEO-R in late 2005.

Partly due to the IP-safe property of the CodeCOOP, upper management at the organization was sufficiently convinced of the CodeCOOP's utility that a third executive decision was made: to deploy the CodeCOOP to the entire engineering organization (on the order of 25 times the size of its research arm). In spring 2006, the CodeCOOP was introduced to the entire company in three ways: (1) 700 posters marketing the CodeCOOP were placed in buildings throughout the company's main campus, (2) an email was sent to an internal distribution list of all engineers company-wide, and (3) an article introducing the CodeCOOP was published in the company-wide newsletter. With this corporation-wide release, the stakeholders at the company now included all engineers and management throughout LEO (not only those employed by LEO-R).

8. SYSTEM ASSESSMENT

Recall that our overarching goal in designing and deploying the CodeCOOP was to build a knowledge-sharing groupware system such that value tensions did not undermine system appropriation. We recognize – as with any real world deployment of a new system as it ramps up – there is no definitive way to assess the overall success of the system and whether values issues were handled well. That said, we provide empirical evidence to help gauge how well the CodeCOOP and our associated Value Sensitive Design processes have worked here. We first examine the viability of the CodeCOOP as a groupware system and then the extent to which value tensions do not cause concern among the user population.

8.1 Participants and Methods

Two types of data were collected: usage data from the CodeCOOP system and semi-structured interviews.

Usage data was collected from the first 40-weeks of system deployment (14 weeks of LEO-R only deployment followed by 26 weeks of LEO-R + LEO deployment) and included information such as the number of (a) IP addresses that had accessed the CodeCOOP, (b) searches conducted, (c) discussion posts, (d) code posts, (e) code downloads, and (f) unsolicited email messages to the system administrator. Since we did not log visitors' identities, we used the number of distinct IP addresses that visited as a proxy for the number of visitors to the system.

Semi-structured interviews lasting roughly 45 minutes were conducted with 18 individuals, 8 CodeCOOP users and 10 CodeCOOP non-users. For CodeCOOP users, interview questions tapped the participants' (1) general impressions of the CodeCOOP (e.g., "What's working for you in regards to the CodeCOOP? What's not working?"), (2) knowledge sharing (e.g., "What's the best piece of information you've gotten from the CodeCOOP?"), (3) value issues around reputation and privacy (e.g., "From your perspective, how well are privacy issues handled in the CodeCOOP? Very well. Moderately well. Not very well. Or just not relevant. What's working for you with the CodeCOOP with regards to privacy? What isn't working?"), (4) distribution lists (e.g., "Do you think it is all right or not all right that the distribution lists were included and are now searchable in the CodeCOOP? Why or why not?"), (5) community (e.g., "Has using the CodeCOOP made you more aware of people and their areas of interest and expertise?"), and (6) managerial policy (e.g., "Do you think managers should consider CodeCOOP contributions in annual evaluations?"). For CodeCOOP non-users, interviews tapped participants' reasons for not using the CodeCOOP. Then non-users were shown screen shots of the CodeCOOP and introduced to the CodeCOOP's privacy and reputation related features. Following this introduction to the system, non-users were asked to evaluate various CodeCOOP technical features and managerial policies. For both users and non-users, structured questions were followed by informal probes to allow participants to elaborate on their perspectives. All interviews were tape-recorded and subsequently analyzed.

8.2 Results

During the course of the 40 weeks, a total of 5,965 distinct IP addresses visited the CodeCOOP. These visitors conducted a total of 30,444 searches, 35,936 views of materials posted within the CodeCOOP, and 2,304 downloads from the code repository. Additional details of the total usage of the system are reported in the Cumulative Total column of Table 4.

8.2.1 Patterns of Usage over Time

For the first fourteen weeks, the system was only available to the relatively small research arm of the company, resulting in very low usage numbers for all types of activities. It was then deployed company-wide starting in Week 15, resulting in an immediate spike in usage. For purposes of this assessment, we focus our attention from Week 15 on. Table 4 reports the average weekly frequency of various types of activities in the system during this initial spike in usage after the CodeCOOP was

Table 4. CodeCOOP Usage Statistics

	Average per Week			
	Wks. 15-18	**Wks. 19-22**	**Wks. 37-40**	**Cum. Total**
Visitors	**831**	**540**	**347**	**5,965**
First time visitors	657	245	112	
Contributions	**19**	**13**	**9**	**436**
Code posts	4	4	1	143
Discussion thread posts	15	9	8	293
Searches	**1232**	**1230**	**1167**	**30,444**
Views	**2193**	**1529**	**1036**	**35,936**
Code post views	1311	712	475	17,677
Archived DL e-mail views	592	655	448	14,204
Discussion thread views	290	162	114	4,055
Code downloads	**69**	**70**	**116**	**2,304**

deployed company-wide (Weeks 15-18), during the four weeks immediately following this initial spike (Weeks 19-22), and during the last four weeks of data collection (Weeks 37-40).

Visitors to the CodeCOOP. The number of visitors to the CodeCOOP increased dramatically after the system was introduced company-wide, reaching a peak of 1,166 distinct visitors during Week 17. Following this initial spike, the number of visitors per week declined steadily but began to stabilize or even increase slightly near the end of the 40 weeks (i.e., within Weeks 37 – 40, detail not shown in Table 4), with an average of 347 distinct visitors per week during Weeks 37 to 40. During these last four weeks, there was an average of 112 first-time visitors to the CodeCOOP each week. Of the 5,965 distinct users of the CodeCOOP, 2,273 (38%) are repeat users who have visited the CodeCOOP on at least two different days and 576 users (10%) have visited the CodeCOOP on at least five different days.

Contributions to the CodeCOOP. Users can contribute material to the CodeCOOP by posting either to the code repository or to a discussion thread. Figure 1 shows the cumulative number of posts of each type over time. Following the company-wide deployment in Week 15, there was steady growth in the number of both types of posts, reaching a cumulative total of 143 code posts and 293 discussion thread posts by Week 40. The flattening of the graph for code posts after Week 30 indicates a slowing in the rate of code posts; in fact, in Weeks 37-40, there was an average of only 1 code post per week.

Searches. The number of searches has remained quite stable at approximately 1200 per week since the company-wide deployment.

Views of Posts. The CodeCOOP contains three types of posts: code posts, discussion thread posts, and archived emails from discussion lists. Following the initial spike in Weeks 15 – 18 the number of views of code posts and discussion thread posts dropped quite sharply, while the number of views of archived DL emails declined more slowly. Toward the end of the 40 weeks, views of all three types of posts stabilized or started to increase again (i.e., within Weeks 37 – 40, not shown in Table 4), with an average of 475 code post views, 448 archived DL email views, and 114 discussion thread views per week during Weeks 37-40.

Code Downloads. The number of downloads of code has been very steady or even increased slightly over time.

Unsolicited Email to CodeCOOP System Administrator. The CodeCOOP provides an email address for submitting feedback (e.g., feature requests, bugs). Approximately 250 unsolicited emails were sent in this manner to the system administrator; of

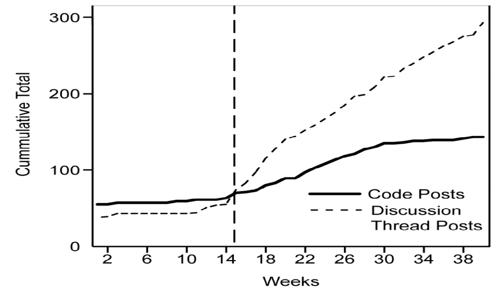

Figure 1. Cumulative number of posts to CodeCOOP

these, only one message mentioned a value tension (requesting an ability to post anonymously to the CodeCOOP).

8.2.2 Interviews

The 10 non-users who were interviewed were asked about their reasons for not using the CodeCOOP. After allowing a period for free response, participants were provided with a list of 19 possible reasons for not using the system and asked to select all that apply. Seven out of 10 selected "I didn't know it existed," and an additional two participants selected "I forgot it was there." Thus 9 out of the 10 non-users specified lack of awareness of the system as a reason for not using it. Apart from one participant who selected a reputation-related issue ("I feel too exposed…"), no other values-related issues were reasons for non-use.

Both users and non-users (after they were introduced to the CodeCOOP) were asked a number of questions about values-related issues. Responses are summarized in Table 5. When asked how well privacy issues were handled in the CodeCOOP, none of the 18 participants (0%) said "not very well." However, 4 participants (22%) indicated that reputation issues were not handled very well. When asked about specific reputation-related issues, 94% of participants said that users should be able to delete their own posts (which is not currently allowed in the system), 83% said users should be able to rank the contributions of others (currently allowed), 39% said users should be able to see who has downloaded their code (not currently allowed), and 33% said that anonymous posts should be permitted (not currently allowed).

With respect to managerial policies, 14 of the 18 participants (78%) said that CodeCOOP contributions should be considered by managers in annual evaluations, but only 6 (33%) thought that there should be a company-wide policy that rewards CodeCOOP usage. LEO-R management is currently considering organizational policy for the CodeCOOP based on these results.

9. DISCUSSION

To what extent can we say that CodeCOOP is a viable groupware system? While what counts as "viable" may differ from situation to situation, reasonable criteria for a groupware system such as the CodeCOOP include (1) a core group of "regular" users, (2) a pool of new users who are still coming to the system, and (3) a useful and growing corpus of knowledge [11, 17].

The usage data indicates that the CodeCOOP is performing moderately well on all three criteria. Since its inception, nearly 6000 distinct users have visited the CodeCOOP. Of these, roughly 10% form a core group of sustained CodeCOOP users who have visited the CodeCOOP on at least five different days. In addition, another roughly 100 new users visit the CodeCOOP each week. Given a large and as yet untapped population of

Table 5. Evaluation Interview Results (N=18 Participants)

	Very Well	Mod. Well	Not Very Well	Not Relevant
Based on your perception of the CodeCOOP, how well are…				
…privacy issues handled?	44%	39%	0%	11%
…reputation issues handled?	11	33	22	17

Should users…	Percent "yes"
…be able to delete their code & discussion posts?	94%
…be able to rank the contributions of others?	83
…be allowed to see who has downloaded their code?	39
… be allowed to post code anonymously?	33
Do you think managers should consider CodeCOOP contributions in annual evaluations?	80
Would you like there to be a company wide policy that rewards use of the CodeCOOP?	33

potential users within the organization, this stream of new users seems likely to continue for some time to come. The usage data indicates that the corpus has been most useful for searches and code downloads as these types of use have sustained the initial burst of usage that occurred when the CodeCOOP was introduced. It appears that viewing the corpus of DL emails, code posts, and discussion threads has been less useful as Table 4 shows an initial burst of usage when the CodeCOOP was first announced and then a sharp decline. However, even for these less popular usage types, Weeks 37 - 40 show an increase in use which could indicate a growing utility for viewing code posts and discussion threads. In terms of a growing corpus, Table 4 and Figure 1 indicate that the corpus is growing both in terms of code and discussion posts contributed (although the rate of code postings has slowed and is an element of concern). Taken together, this usage data suggests the CodeCOOP is alive and establishing itself within the organizational culture.

Moreover, in their interviews participants identified unique knowledge sharing and community functions provided by the CodeCOOP that set it apart from other available technologies. For example, in contrast to the Internet, the CodeCOOP provides access to cutting edge internal technical information (e.g., "for sort of the newer things that are internal only, that sort of thing … new technologies.") as well as to IP-safe source code (e.g., "It's probably about the best start the company has right now for doing source-code sharing that does relate to the software publishing business purpose"). It also supports a sense of community by increasing awareness among community members and providing easy access to otherwise fragmented community dialog (e.g., "being able to track that sort of shared comments over time … I think it's a useful community service").

Having established a reasonable degree of viability of the CodeCOOP, we turn next to value-oriented aspects of the system. By and large, the value tensions we identified early in our initial design process – primarily those of privacy and awareness – and sought to minimize through our design did not appear to be problematic. We bring two sorts of evidence to bear. First, the absence of references to value issues both from CodeCOOP users in the form of unsolicited emails received by the system administrator (1 message out of roughly 250 messages) and from CodeCOOP non-users in the form of the reasons they provided for their non-use (virtually no value-related issues were cited for non-use though several such issues were among the 19 reasons listed which participants selected among). Second, when asked explicitly how well the CodeCOOP handled privacy issues, none of the interview participants (0%) said "not very well." Taken together, these results suggest that privacy issues were satisfactorily addressed in the CodeCOOP design and, more generally, provide some support for the proposition that the Value Sensitive Design methods employed here were successful in mitigating at least some potential value tensions.

That said, at least one value issue emerged more strongly as our design progressed: that of reputation. In our initial pre-development assessment, only 1 of 6 questions on reputation was classified as a value dam while both reputation-related benefits were identified as flows (refer back to Table 2). Thus, at that time, we did not give a great deal of attention to reputation in the design. However, at the time of this evaluation 4 participants (22%) indicated that reputation issues were not handled very well in the CodeCOOP. To remedy this situation and as suggested by the interview data, we expect to implement a delete feature that would allow users to remove their own code and discussion posts.

10. CONTRIBUTIONS AND FUTURE WORK

Overall, integrating Value Sensitive Design practices with the organization's design processes was relatively seamless. We experienced a good deal of success applying methods developed in a research environment to the industry setting, including conducting stakeholder analyses, harms and benefits analyses, and an empirical assessment of stakeholder views and values. Though we initially expected our methods to be too slow and cumbersome for the industry context, we did not find that to be the case. The industry context also pushed us to identify a gap in established methods and to develop a new method in response. Thus, the work reported here makes four key contributions:

- Provides a "proof-of-concept" that established Value Sensitive Design principles and methods developed in a research environment can be used to good effect in an industry setting. Moreover, these Value Sensitive Design methods can be employed alongside of other design and usability practices.
- Development of a new design method for addressing value tensions in design, Value Dams and Flows. This method is intended to be light-weight, tangible, and readily integrated with other design practices.
- Demonstrates the co-evolution of technology and organizational policy. Specifically, using Value Dams and Flows we both designed technical features and suggested managerial policies to support system use.
- Documents emergent stakeholders and value issues in the on-going design and deployment of a groupware system and provides a value-sensitive analysis.

In turn, the work reported here also points to future directions:

Refinement and Validation of the Value Dams and Flows Method. Further research is needed on how to best determine threshold values for value dams and flows in a project. Ideally the value dams threshold would account for how strongly concerns are held, severity of the potential harm, and who might be harmed; similar dimensions of benefits are likely relevant for value flows. In addition, while our experience suggests the value dams and flows method successfully guided our design and supported appropriation, additional case studies and perhaps experimental designs could be employed to refine and validate this method.

Increase Involvement of Indirect Stakeholders. Due to organizational constraints, we were unable to involve some categories of indirect stakeholders from both within (e.g., high-level managers) and beyond (e.g., customers or stockholders) the organization. Finding methods to incorporate views and values of indirect stakeholders of technology built in industry remains future work.

Broader Spectrum of Organizational Policies and Practices. In this project, we considered managerial policies to support contributors of the CodeCOOP. However, future work would do well to explore a broader range of policies and practices that organizations may employ to support values around technology.

In conclusion, values matter to people who use and are affected by the use of information systems – so much so that at times unresolved value tensions have contributed to the failure of information system appropriation [11, 14]. If designers and developers in fast-paced, bottom-line oriented industry settings are to account for human values, we need to provide light-weight, principled methods for them to do so. Toward that end, this work endeavors to develop design methods that researchers and designers, in academia and industry, can use to help account for human values in their design process – both for groupware and other types of information systems.

11. ACKNOWLEDGMENTS

We would like to thank those who contributed in various ways to the project's success: A.J. Bernheim Brush, Alan Borning, Danyel Fisher, Nathan Freier, Darren Gehring, Jonathan Grudin, Rob Gruen, Peter Kahn, John Lin, David McDonald, Grace Preyapongpisan, Rachel Severson, Marc Smith, and Gina Venolia. This material is based, in part, upon work supported by the National Science Foundation under Grant No. 0325035. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

12. REFERENCES

[1] Bellotti, V. and Sellen, A. Design for privacy in ubiquitous computing environments. In *Proc. of ECSCW* (Milan, Italy, Sept. 13-17, 1993). Kluwer Academic Publishers, Dordrecht, The Netherlands, 1993, 77-86.

[2] Boyle, M., Edwards, C., and Greenberg, S. The effects of filtered video on awareness privacy. In *Proc. of the 2000 ACM Conf. on CSCW* (Philadelphia, PA, USA, Dec. 2-6, 2000). ACM, N.Y., N.Y., 2000, 1-10.

[3] Flanagan, M., Howe, D. C., and Nissenbaum, H. Values at play: design tradeoffs in socially oriented game design. In *Proc. of CHI* (Portland, OR, USA, April 2-7, 2005). ACM, N.Y., N.Y., 2005, 751-760.

[4] Friedman, B., Kahn, P. H., Jr., and Borning, A. Value Sensitive Design and information systems. In *Human-Computer Interaction and Management Information Systems: Foundations*. M.E. Sharpe, Armonk, N.Y., 2006, 348-372.

[5] Friedman, B., Lin, P., and Miller, J. K. Informed consent by design. In *Security and Usability: Designing Secure Systems That People Can Use*. O'Reilly, Cambridge, 2005, 495-521.

[6] Friedman, B. and Nissenbaum, H. Bias in computer systems. *ACM Trans. on Inf. Syst.*, 14, 3 (July 1996), 330-347.

[7] Friedman, B., Smith, I. E., Kahn, P. H., Jr., Consolvo, S., and Selawski, J. Development of a privacy addendum for open source licenses: Value Sensitive Design in industry. In *Proc. of UbiComp 2006* (Orange County, CA, USA, Sept. 17-21, 2006). Springer-Verlag, Berlin, 2006, 194-211.

[8] Gaw, S., Felten, E. W., and Fernandez-Kelly, P. Secrecy, flagging, and paranoia: adoption criteria in encrypted e-mail. In *Proc. of CHI* (Montréal, Québec, Canada, Apr. 22 - 27, 2006). ACM, N.Y., N.Y., 2006, 591-600.

[9] Greenbaum, J. M. and Kyng, M. *Design at Work: Cooperative Design of Computer Systems*. Lawrence Erlbaum Associates, Hillsdale, N.J., 1991.

[10] Grønbæk, K., Grudin, J., Bødker, S., and Bannon, L. Achieving cooperative system design: shifting from a product to a process focus. In *Participatory Design: Principles and Practices*. Lawrence Erlbaum Associates, Mahwah, N.J., 1993, 79-96.

[11] Grudin, J. Groupware and social dynamics: eight challenges for developers. *Comm. of ACM*, 37, 1 (January 1994), 92-105.

[12] Jancke, G., Venolia, G. D., Grudin, J., Cadiz, J., and Gupta, A. Linking public spaces: technical and social issues. In *Proc. of CHI* (Seattle, WA, USA, Mar. 31 - Apr. 5, 2001). ACM, N.Y., N.Y, 2001, 530-537.

[13] Nissenbaum, H. Will security enhance trust online, or supplant it?. In *Trust and Distrust Within Organizations*. Russell Sage Foundation, N.Y., N.Y., 2004, 155-188.

[14] Orlikowski, W. J. Learning from Notes: Organizational Issues in Groupware Implementation. In *Proc. of the 1992 ACM Conf. on CSCW* (Toronto, Canada, Oct. 3 - Nov. 4, 1992). N.Y., N.Y.: ACM, 1992.

[15] Palen, L. Social, individual, and technological issues for groupware calendar systems. In *Proc. of CHI* (Pittsburgh, PA, USA, May 15-20, 1999). ACM, N.Y., N.Y., 1999, 17-24.

[16] Palen, L. and Dourish, P. Unpacking "privacy" for a networked world. In *Proc. of CHI* (Ft. Lauderdale, FL, USA, Apr. 5-10, 2003). ACM, N.Y., N.Y., 2003, 129-136.

[17] Preece, J. and Maloney-Krichmar, D. *Online Communities: Focusing on Sociability and Usability*. Lawrence Erlbaum Associates, Mahwah, N.J., 2003.

[18] Van House, N. A., Butler, M. H., and Schiff, L. R. Cooperative knowledge work and practices of trust: sharing environmental planning data sets. In *Proc. of the 1998 ACM Conf. on CSCW* (Seattle, WA, USA, Nov. 14-18, 1998). ACM, N.Y., N.Y., 1998, 335-343.

[19] Zheng, J., Veinott, E., Bos, N., Olson, J. S., and Olson, G. M. Trust without touch: jumpstarting long-distance trust with initial social activities. In *Proc. of CHI* (Minneapolis, MN, USA, April 20-25, 2002). ACM, N.Y., N.Y., 2002, 141-146.