European Journal of Information Systems (2006) 15, 473–485
© 2006 Operational Research Society Ltd. All rights reserved 0960-085X/06 $30.00
www.palgrave-journals.com/ejis

The incompatibility of knowledge regimes: consequences of the material world for cross-domain work

Jennifer A. Howard-Grenville[1] and Paul R. Carlile[2]

[1]Organizational Behavior Department, Boston University School of Management, Boston, MA, U.S.A.; [2]Information Systems Department, Boston University School of Management, Boston, MA, U.S.A.

Correspondence:
Jennifer A Howard-Grenville, Organizational Behavior Department, Boston University School of Management, 595 Commonwealth Avenue, Boston, MA 02215, USA.
Tel: +1 617 358 2279;
Fax: +1 617 353 5244;
E-mail: jahg@bu.edu

Received: 3 August 2005
Revised: 23 October 2005
Accepted: 24 August 2006

Abstract

In this paper, we argue that successful integration of knowledge across work domains in the short-term can mask the generation of long-term consequences. We explore a setting, the introduction of environmental considerations into semiconductor manufacturing, where the eventual adoption of common measurement artifacts and associated practices enabled knowledge integration, but failed to address significant underlying consequences. Drawing from observational, interview, and archival data we develop an understanding of the work practices of the Tech and EnviroTech groups as structured by the material world and broader collective conventions. We introduce the concept of knowledge regime to outline the differences in knowledge across these work domains. More specifically, we find that differences in the causal specificity and developmental time horizon of knowledge and the measurement artifacts that result contribute to the relative power of one knowledge regime over another. Understanding these sources of incompatibility provides insight into the design requirements of information systems as boundary objects for knowledge integration, but also specifies the potential limits to any design effort.

European Journal of Information Systems (2006) **15,** 473–485.
doi:10.1057/palgrave.ejis.3000635

Keywords: knowledge integration; knowledge regimes; material world; organizational boundaries

The challenges of integrating knowledge across the boundaries between specialized domains are now well documented (Brown & Duguid, 2001; Carlile, 2002). Specific roles for actors as boundary spanners (Allen, 1977; Hargadon & Sutton, 1997), artifacts as boundary objects (Star, 1989; Karsten *et al*, 2001; Carlile, 2002) and processes as boundary practices (Brown & Duguid, 2001; Merali, 2002; Carlile, 2004) have been identified as helpful in meeting these challenges. In this and other research, however, there is an unexplored assumption that the benefits of integrating knowledge across boundaries outweigh the costs generated. Indeed, in cases where groups are united around a common goal and can find common ways to express their interests and understand those of others (Bechky, 2003) this assumption may be true. However, in many cases groups simultaneously hold multiple interests that may or may not support each other or broader organizational goals (Howard-Grenville, 2005). This is a reminder that differences in knowledge at the boundary between groups are not always convergent, and can be and sometimes should be seen as divergent (Gheraradi & Nicolini, 2002; Carlile, 2004).

In this paper, we argue that in some situations knowledge is incompatible and further that under these circumstances integration in the near term can mask negative consequences in the long-term. This incompatibility goes much deeper than just differences in work practices, to differences in what we call *knowledge regimes*. We define knowledge regimes as the nested connections between the material realities engaged by work practices, the work practices themselves, and the larger collective conventions that reflect and account for the appropriate use of such practices. In recognizing the material world as a potentially important starting point for the shaping of work practice, the concept of knowledge regimes resembles Thévenot's (2001) use of the term 'pragmatic regimes' to represent the 'relevant realities' in which practices are embedded and the social devices that govern our ways of engaging with and articulating the world around us. We examine differences between knowledge regimes to understand how these differences might generate long-term consequences for knowledge integration, specifically in the domain of technological innovation within firms.

We focus empirically on the semiconductor manufacturing industry, the producers of computer and other electronic 'chips.' What makes this industry of particular value in elaborating the concept of knowledge regimes is the fact that innovation in the industry derives from fundamental material properties of silicon and has continued at a relentless pace for four decades. Industry founder Gordon Moore observed in 1965 that the number of electronic components on a chip would roughly double every 2 years. Now referred to as Moore's Law, this observation has shaped industry forecasts and expectations that influence the actions of a large set of global actors, including semiconductor firms, suppliers, competing firms, scientific researchers and even end consumers. These expectations, and the investments they have motivated, have in turn enabled the continued exploitation of the material properties of silicon, suggesting a recursive relationship between the social and material. Nonetheless, this relationship would be impossible to maintain in the absence of particular material properties.

Within this overarching regime, symbolically represented by Moore's Law, we explored the work practices of members of technical engineering ('Tech') and environmental engineering ('EnviroTech') groups at 'Chipco' (a pseudonym), a major semiconductor manufacturer. Observing how members of Chipco sought to mitigate environmental impacts of its manufacturing processes, we found that, over time, complex new environmental issues were increasingly represented as manufacturing constraints. Accounting for the success of one project, an engineer noted 'it was the first time we treated an [environmental device] like process equipment.' However, such technical fixes could be far from environmentally optimal, and in this case shifted the waste from the air to water, and subsequently to landfill. While they actively sought to mimic the practices of Tech, members of EnviroTech nonetheless recognized that such approaches failed to address the 'cost and complexity of treating the 'lifecycle'' of hazardous materials.

In addressing this puzzle – why a seemingly successful integration effort had nonetheless left the environmental issues unsolved – we inductively analyzed data derived from observation, interviews and archival sources to identify how knowledge was structured at the material, work practice, and broader conventional level for the two groups. We found considerable differences between the knowledge regimes themselves and uncovered characteristics that explain differences in their persistence and relative power. The semiconductor manufacturing regime is persistent and powerful because it produces knowledge that tightly specifies relationships between cause and effect, is developed on a short time horizon, and is represented through measurement artifacts (i.e., Moore's Law and all its implications) that are relatively clear and widely shared. By comparison the environmental knowledge regime produces knowledge that is much less causally specific, developed over a much longer time frame, and has very few unambiguous and uncontested measurement artifacts to deploy among the larger set of actors involved. With less predictable and slower to accumulate knowledge, the environmental knowledge regime was subject to a 'pull' of its work practices and conventions by the more powerful semiconductor manufacturing regime.

These key characteristics – causal specificity, time horizon, and metrics – can be generalized to other regimes and used to understand the relative power of one knowledge regime over another and the potential costs associated when aligning work practices across them. The challenges and paradoxes of working across boundaries in the development of information systems are well documented (Huang, Newell, & Pan, 2001; Karsten *et al*, 2001; Merali, 2002), so understanding the characteristics of particular knowledge regimes is important when considering the design of information systems and their successful deployment. Although this study is not focused on a particular information system, exploring the sources of incompatibility will help us better understand the design requirements of information systems and other artifacts as boundary objects as well better recognize the potential limits of such design efforts.

Further, advances in semiconductor manufacturing have direct implications for the capabilities of information technologies; Moore's Law has enabled the power of information systems to increase at a rapid pace and generates expectations that this pace of innovation will continue. This paper calls attention to the properties that underlie such expectations, and the potential difficulty that they pose for representing and acting on knowledge that accumulates more slowly and is grounded in different material realities. Where other analyses of information system deployments have focused on the role of power and language in shaping opportunity and action (McGrath, 2002), this paper brings attention to another level, the material world, and explores how it can, in turn, influence patterns of practice, discourse, and power.

Knowledge integration and knowledge regimes

Integrating knowledge across organizational boundaries is widely recognized as challenging especially when it requires the coordination of multiple groups with specialized expertise, as in product or process development, or technological innovation (Dougherty, 1992, 2001; Leonard-Barton, 1992). While such boundaries can be overcome by the 'externalization' of tacit, situated knowledge (Nonaka, 1994) through close, frequent interactions between groups (Hansen, 1999), co-location (Tyre & von Hippel, 1997), the use of material boundary objects for joint problem-solving (Carlile, 2002; Bechky, 2003), or the work of individual knowledge brokers who translate meaning (Hargadon & Sutton, 1997), these approaches do not typically go beyond specific episodes of integration to consider whether there are limits to integration, or to identify the costs of integrating knowledge across domains. The majority of the literature on knowledge integration focuses on a variety of settings such as project based organizations (Hansen, 1999; Scarbrough *et al*, 2004), product development (Hargadon & Sutton, 1997; Carlile, 2002) or technology implementation (Huang *et al*, 2003; Shell & Scholz, 2003) – none of which cast a large enough net to examine the long-term consequences of seemingly successful integration. Further, practice-based approaches almost exclusively focus on the challenges and costs associated with moving knowledge at only the level of work practice (Suchman, 1987; Orlikowski, 2002; Carlile, 2002), and so avoid paying attention to other constraints such as material and institutional influences (Nelsen & Barley, 1997; Dougherty, 2001).

While some recent work acknowledges that knowledge integration may be costly (see Zollo & Winter, 2002) there is very little understanding of the mechanisms that give rise to the costs. Indeed, studies of boundary practices have typically focused on stable product development settings, where a common product or costumer mitigates some of the differences that are present when different goals are desired, or technical and non-technical understandings must be negotiated (Bechky, 2003). In settings characterized by divergent goals, as in the integration of new, environmental considerations into ongoing manufacturing process development, we need a way of understanding differences at boundaries as embedded in material realities and broader conventions that are themselves possibly divergent.

We build upon Thévenot's (2001) use of the term 'pragmatic regimes' in generating our concept of knowledge regimes. The concept of a regime has a history within the social sciences from its initial use by Marx to a more recent expression by Foucault (i.e., regime of power; Foucault, 1980) to indicate the complex forces that shape social action. Thévenot offers a practice based (Bourdieu, 1977) and grounded expression of pragmatic regimes because he pays attention to the *material world* as a means of anchoring practice (Boltanski & Thévenot, 1999). Concerned with work that conceptualizes the social primarily as 'starting from common frames of understanding,' Thévenot provides an alternative by arguing that pragmatic engagements of individuals with the world are a starting point for understanding action. Pragmatic engagements become conventionalized only to the extent that they are publicly evaluated and judged as appropriate engagements. Thus, Thévenot's conception of a pragmatic regime refers first to the material realities in which practices are embedded, and secondarily to the social devices that govern engagement with and articulation of these realities.

Reflecting Thevenot's concerns, we define knowledge regimes as the nested connections between the material reality engaged by work practices, the work practices themselves, and the larger collective conventions that reflect and account for the appropriate use of such practices. This anchors practices in the material world and allows material realities to produce initial contours for the work practices. It also draws attention to broader conventions that reinforce and magnify the connections between materiality and practice. This anchoring in materiality and a multi-level approach provides a unique basis from which to describe knowing in practice. Seeing knowledge regimes as comprised of connections between materiality, practices, and collective conventions also provides a way of understanding the relative power of one regime over another and identifying sources of incompatibility that may not be evident in an examination of work practices alone.

Setting and methods

We used qualitative data sources, including ethnographic fieldnotes, interview transcripts, and archival data to explore the embeddedness of work practices at Chipco in two different knowledge regimes and the consequences of this. We approached our data collection and analysis on three levels: the material; work practice; and industry convention levels. We considered the structure of two knowledge regimes, one surrounding semiconductor manufacturing work, and one surrounding environmental work.

The primary method used to collect data for the material and work practice levels was an ethnographic field study undertaken by one author, who was a full-time (45 + hours per week) participant observer at Chipco over the course of a nine-month period. Chipco is one of the world's largest manufacturers of microprocessor 'chips' used in computers. Headquartered in the U.S., Chipco's semiconductor manufacturing facilities ('fabs') are located in eight locations worldwide. At the time of this study, Chipco employed 68,000 people, 75% of whom were involved in chip manufacturing, including manufacturing development and support.

New manufacturing process development is a prominent activity within Chipco, as a major new manufacturing process generation is introduced every two years to enable the production of faster, more powerful chips. Roughly one-third of the hundreds of process steps required to manufacture a chip undergo a major change (e.g., new equipment) and many of the remaining steps undergo some significant change (e.g., same equipment but new operating parameters, new chemicals or gases, etc.) on this 2-year cycle. Discrete projects to develop or improve particular process steps within this overall development effort typically span 6 months to 1 year, and involve testing, modifying, and optimizing equipment and procedures at a dedicated development fab. A typical process development project at Chipco is led by several members of the 1,500-person process technology development group, 'Tech,' with managers and engineers from other groups involved as needed.

Starting in the early 1990s and in response to a handful of incidents in which environmental regulations or concerns threatened to hold up the implementation of new manufacturing processes, Chipco had made several formal changes to help integrate environmental considerations into the development of its manufacturing processes. These eventually led to the establishment of a new 11-member group, EnviroTech, created to work with Tech to mitigate environmental impacts of process technologies at the design and development stage. Over 8 years, EnviroTech (and its precursors) had become increasingly sophisticated and successful in how they incorporated the environmental considerations into Tech's process development projects. As noted earlier, they worked to represent the environmental issues themselves using measurements valued in Tech, for example, showing how failure to address an environmental issue would result in limited manufacturing capacity. Earlier efforts had drawn more explicit attention to the environmental considerations by, for example, calling attention to emissions levels or community concern, but had not generated action within Tech. EnviroTech had increasingly mimicked the work practices prevalent within Tech, coming to rely more heavily on providing data to drive decision making and orienting their work to technical solutions that could be developed and deployed rapidly.

Data collection

As a participant observer within the EnviroTech group, the researcher entered the field with a goal of understanding how environmental issues were being surfaced, articulated, and acted upon within Chipco's process development activities and she wrote daily fieldnotes to capture her observations. She also collected data through semi-structured interviews performed and documents collected during the course of observation. We used this data primarily to inform our understanding of the work practices within Tech and EnviroTech. We developed an understanding of the material aspects of semiconductor manufacturing through the participant observation and augmented our understanding of it through the use of technical sources on the semiconductor manufacturing process (Van Zant, 1997; Murphy *et al*, 2003).

At the industry level, we collected all versions of the National Technology Roadmap for Semiconductors (NTRS) published over the last 11 years. First published in 1992, we were able to obtain copies of all the subsequently published Roadmaps: 1994, 1997, 1999, 2001 and 2003. The documents are more than 200 pages in length and represent industry consensus on process technology needs for 15 years in the future. Coordinated by the Semiconductor Industry Association (SIA) the goal of the Roadmap documents is to provide a 'framework for guiding R&D … to meet the increasingly complex technology needs of the semiconductor industry' (SIA, 1994: p ix). From 1999 onwards, the Roadmaps have been international efforts (now called the International Technology Roadmap for Semiconductors, ITRS) and have been jointly sponsored by the U.S., European, Japanese, Taiwanese, and Korean semiconductor industry associations. Sections of the Roadmap are prepared by representatives from semiconductor manufacturers, suppliers, research consortia, academia, and government. As such, they capture the state of thinking on the critical research needs and future directions of the industry for each year they are published, and they reflect the evolution of this thought over time.

Data analysis

Analysis of the Chipco data was performed in two steps. First, the participant observer constructed a detailed description of the Chipco culture, and the Tech and EnviroTech subcultures in particular (see Howard-Grenville, 2006), by coding for emergent themes (Glaser & Strauss, 1967; Miles & Huberman, 1994). In a second step, she analyzed in detail seven projects that had been initiated over a 6-year time span and coded for practices that were used by EnviroTech in interfacing with Tech. The first analysis enabled us to understand the work practices in Tech and their relationship to the material reality of semiconductor manufacturing, while the second analysis enabled us to see shifts over time in EnviroTech's work practices.

We then analyzed the national/international Roadmap documents, focusing primarily on three sections, the executive summary, the section on lithography, and the section on environment, safety and health (ESH). Within the executive summaries of each Roadmap document we noted the stated goals of the Roadmap and key challenges facing the industry. We used the lithography section as representative of how core semiconductor manufacturing technologies are treated in the Roadmap, consistent with the industry's own interpretation of lithography processes as driving all other aspects of the chip manufacturing process (SIA, 2001). We compared and contrasted the format and content of the tables in the lithography section with those in the ESH section, and compared the text of each section. We were concerned with discerning similarities and differences in what information was displayed, how it was displayed, and how uncertainties and key challenges were represented.

The knowledge regime of semiconductor manufacturing

We describe here how the semiconductor manufacturing knowledge regime is produced and reproduced at the material, work practice and industry convention levels. Two key native characteristics – 'scalability' and 'divisibility' – emerged from our analysis as central to this regime. Table 1 summarizes these characteristics and how they are manifest at each of the three levels. The arrows shown pointing up from the bottom of the table depict the fact that these critical characteristics of the knowledge regime are anchored at the material level but are also reinforced and magnified at the other levels.

Material level

Today's chips are manufactured using techniques that were first developed in the 1950s. The basic premise of chip manufacturing takes advantage of the natural semiconducting properties of silicon to produce integrated circuits (electronic components like transistors, resistors and capacitors, and electrical connections between them) on a thin silicon 'wafer.' By masking and patterning certain areas of the wafer, depositing material or 'growing' new material by altering properties of the silicon itself, the components of the integrated circuits are built (Van Zant, 1997). Critically, a region of any size can be electrically altered to make an electronic component of virtually any size on very high-purity silicon wafers. Effectively the only limit to *scaling* – the reduction in the physical size of the components, or 'features,' on a chip – lies in the limits on the precision of the technologies used for four core operations (patterning, etching, deposition and implanting). As these technologies have improved the minimum feature sizes on chips has decreased by a factor of about 100 over 40 years.

The material property of scalability has two key implications for chip manufacturing. First, more and more components can be packed into the same chip area, making the chips themselves more powerful and faster. Second, the cost of producing chips has *not* historically risen at the same rate as the component density, making chips progressively cheaper relative to their functionality. Few other technology platforms demonstrate such radical scaling over such a huge range. Indeed, industry founder Gordon Moore liked to draw the analogy that 'if the automobile industry was like the (scalable) semiconductor industry, a Rolls Royce would get half a million miles to the gallon and it would be cheaper to throw it away than to park it.'

A second fundamental feature of chip manufacturing is *divisibility*. The depth, not just the area, of the silicon altered in the semiconductor manufacturing process can be precisely controlled. Material can be deposited or grown in layers as small as a few nanometers (one-billionth of a meter). This allows for the sequential assembly of a chip, layer by layer. While at least one hundred process steps are needed to manufacture a given chip, they are repeated sequences of eight basic unit

Table 1 Key aspects of semiconductor manufacturing knowledge regime

Industry convention level	↑	• Industry roadmaps identify technology trends and critical technology needs for 15 years into future • Roadmaps represent consensus between manufacturers, suppliers, researchers, and others.	↑	• Roadmaps defined for each major process area, and further broken down into particular process steps • Dedicated suppliers for process equipment by major process step
Work practice level	SCALABILITY	• Work practices oriented around specific improvements in process parameters to shrink components by given scale • Past success at scaling gives confidence in ability to overcome future technical challenges	DIVISIBILITY	• Roles specialized by chip layer and process step • Optimization of each process step, with integration achieved by dedicated specialists
Material level		• Electronic components on chip physically shrunk, resulting in faster, higher performance chip		• Chips built up layer by layer using discrete processing steps and equipment • Minutely thin top layer of silicon (only a few nm deep) is altered sequentially

operations (Murphy *et al*, 2003), each of which is performed using specific, discrete process equipment.

Work practice level

Key aspects of the organization of work and members' understanding of the work at Chipco reflect the importance of the scalability and divisibility of the underlying physical processes. First, scalability affords the development of precise and predictable goals for individuals and work groups. Work in Tech revolved around setting and achieving these specific goals, labeled 'target specifications,' for discrete manufacturing process steps. Tech managers established these specifications for each of the process steps under development, inevitably by scaling them from the last manufacturing generation. In fact, one engineer joked that managers' calculators were capable of multiplying by only one number, 0.7, which is the factor by which linear dimensions must shrink to double the density of components on a chip. As a result of scalability these specifications were central to the work and culture of Tech engineers. One suggested that 'there are three things that count in the life of a [Tech] engineer: what your group leader tells you to do, what the goals for your process module are, and what your performance review says.'

The divisibility of the manufacturing process was reflected in the formal and informal organization of work, with roles highly specialized around discrete types of processes and even discrete chip layers. Interaction was intentionally limited, and one Tech engineer observed that he and his peers 'all have their blinders on' so they could focus exclusively on their part of the process. An 'etch' engineer working on a particular chip layer may have very little interaction with etch engineers working on other layers, or with a 'polish' engineer working on the same chip layer. Interaction between managers was similarly minimized – attendance at meetings to decide on individual process steps and consider integration between them was strictly limited to those who had a 'need to know.' Such meetings, one manager noted, were 'never just FYI.' This clear sequential division of labor at the work practice level was enabled by the material divisibility of semiconductor manufacturing itself.

Together scalability and divisibility allowed for a unique organization of work practices. They enabled the interfaces between the specialized domains required for chip manufacturing to be well defined and for future dependencies between these domains to be anticipated and managed. Further, within and across each specialized area, they enabled relatively simple and specific targets to be set, and progress against these targets to be assessed.

Industry/collective convention level

Finally, our review of the industry's technology Roadmaps shows that scalability and divisibility are also prominent characteristics of knowledge at the industry level. First, the Roadmaps identify 'technology nodes' that mark 'the achievement of significant advancement in the process technology' (SIA, 2003). Each technology node is defined in terms of a metric that represents the minimum feature size to be achieved on a chip (expressed, prior to 2001 as the wavelength of light needed to pattern such a feature size, and, after 2001 as the 'half-pitch' or physical size patterned). Each Roadmap contains multiple individual tables that show technology trends for specific process steps and each table is delineated by both the common technology nodes and the future year running across the top two rows. At the level of industry expectations, time and technology advances appear to enter into lock step, reflecting the physics of scalability.

Over the last decade the semiconductor manufacturing process has become somewhat more complex, with concerns that fundamental physical limits to scalability will be reached (SIA, 2003). Despite this, the Roadmaps continue to reflect a strong adherence to the assumptions of scalability; indeed, as the technical challenges to achieve scalability mount, they become specified in greater detail. While the 1994 roadmap contained 27 technology trend tables for specific aspects of the manufacturing process, by 1999 there were 96 such tables, and by 2003 there were 124. Each of these is delineated by time (year) and technology node, and each almost exclusively depicts the numeric scaling of various process parameters over a 15-year time horizon.

Furthermore, the Roadmaps reflect the divisibility of the manufacturing process in their division into separate chapters for separate groups of major process steps. The chapters are delineated virtually identically across all of the years covered by the Roadmaps, suggesting the stability of the types of processes that provide the core of manufacturing and an ongoing ability to separate them from each other. A chapter in each roadmap is dedicated to each set of major manufacturing process steps, with separate chapters dedicated to 'cross-cutting' issues that affect multiple stages in the manufacturing process (including factory integration, environment, safety and health, yield enhancement, and modeling and simulation). Integration within the Roadmap document is achieved primarily by the fact that all of the technology trend tables, including those whose contents have no physical connection to scalability, are united by the technology node that delineates progressions through time. This is also an indicator of the strong and necessary connection between scalability and divisibility.

Roadmaps are artifacts used by the entire industry to define the interfaces among a large set of organizational actors, including those who produce computers and other products using the industry's chips. The fundamental assumption of scalability guides this large number of actors as they define future needs and direct resources in accordance with adherence to the common convention of Moore's Law. The assumption of continued scalability is not merely arbitrary but is grounded in the physical facts of the materials and manufacturing techniques used.

Persistence of the regime

Scalability and divisibility across all levels together produce a very robust knowledge regime for semiconductor manufacturing. Moore's Law is a simple yet powerful indicator of the persistence of this knowledge regime that was spawned over 40 years ago. Its effects are felt widely through the entire IT industry and even to consumers who now expect and are entrained to major improvements in the power and functionality of computers on a 2-year cycle. The relative power of the semiconductor manufacturing knowledge regime will become clearer when we compare it to the environmental knowledge regime.

The environmental knowledge regime

The knowledge used by those attempting to introduce environmental considerations at Chipco was also grounded in a material reality and reflected in broader conventions. However, unlike semiconductor manufacturing, this knowledge regime was not characterized by scalability and divisibility across all levels. Indeed, uncertainty, complexity, and the inclusion of the expectations of multiple diverse actors better characterized this regime. Key features of each level of the two regimes are summarized in Table 2.

Material level

At the material level, environmental work associated with semiconductor manufacturing was typically neither scalable nor divisible. In fact, the material reality engaged by environmental specialists appeared to operate in quite a different way. Where scalability afforded predictability, the natural environment seemed, for many at Chipco, to be anything but predictable. For example, one Tech manager derided EnviroTech managers for not having forecast and prepared for the discovery that a critical process gas was a potent greenhouse gas. The issue, she noted, 'had just come along from somewhere.' However, even in scientific circles, the gas had only been recently identified as a greenhouse gas, and the quantity released to the atmosphere and its global warming potential (and key metric established by scientific consensus) remained uncertain.

One physical characteristic of the natural environment is that the data needed to understand the interactions between manmade materials and different aspects of the ecosystem are extremely complex. Furthermore, effects often take place distant in time and place from their origins. For example, greenhouse gases are emitted locally, but are expected to have global effects that may only show up fully some decades in the future. Deciding which greenhouse gases are most potent, how to limit their emissions, and who should be responsible for what portion of emissions reduction moves the scientific uncertainty associated with understanding the gases and their climatic effects into the realm of public policy, vastly increasing the number of voices that must be accommodated. In such a setting, environmental specialists experience much less control over the material world than do their Tech counterparts. They cannot easily define and run the experiments that will give them an answer. Furthermore, their findings in one area or for one issue cannot be simply scaled as in the linear and controllable regime of semiconductor manufacturing.

Certain things *are* well known in environmental science, among them that environmental problems are not easily divisible. Materials do not simply go away, but they get transferred between media – from air, to water, to land, and vice versa – or assume different forms that may have different effects. As one EnviroTech engineer noted, this phenomenon posed a dilemma for controlling two different environmental impacts of chip manufacturing:

> 'Because you need [chemical F] to clean [compound S], you are operating within a box; with [chemical F] you are either going to get lots of [toxic air emissions] and a little [greenhouse gases], or lots of [greenhouse gases] and a few [toxic air emissions] – [chemical F] has to come out one way or another.'

Table 2 Comparison of semiconductor manufacturing and environmental knowledge regimes

	Semiconductor manufacturing	Environmental
Industry convention level	• Broad consensus on detailed direction and timing of technology development • Adherence to and reproduction of Moore's Law	• Broad principles (e.g. precautionary principle) but no single one adopted widely by a broad set of actors • Political and social expectations matter, and are subject to change
Work practice level	• Work targets oriented around scalability • Roles specialized by chip layer and process step • Minimal, routinized integration	• Increased attention to use of 'data' • Mimicking Tech's focus and decision processes
Material level	• Scalability and divisibility of manufacturing techniques due to properties of silicon	• Complexity and uncertainty of environmental impacts • Connections and tradeoffs between impacts to different media (air, land, water); not divisible

The understanding of 'operating within a box' suggests that the material world relevant to the environmental specialists presents constraints, as it does to those within Tech. But the constraints are almost never predictable, and past successes do not scale to enable future ones.

Work practice level

At the level of work practices, EnviroTech increasingly shifted toward the use of technical language and metrics that mimicked those of Tech. This is well illustrated by the 'Blue Skies' project. Early in the project EnviroTech managers argued that visible, but harmless emission plumes were 'emotive' for certain communities and suggested that 'the public affairs people don't want to have to explain it anymore.' When this approach failed to win converts among Tech and other groups, they adjusted their language and metrics and crisply presented data on chemical concentrations, scrubber efficiency, and exhaust dilution, and the (albeit uncertain) impact of this on permitting efficiency for future manufacturing facilities. Intentionally manipulating the message in this way framed the issues as valid considerations for Tech, and won them approval and support for the project.

EnviroTech explicitly sought to adopt Tech's approaches. Like Tech, it set specifications, assigned target deliverables, gathered data, performed experiments, and obtained results. However, the organization of the group and their roles did not reflect the divisibility and scalability of the work as in Tech. Indeed, as members of a small group, EnviroTech managers and engineers saw themselves as wearing many hats. One manager who was frustrated by Tech's insistence that he tell them a clear target for greenhouse gas emissions spoke of the complex and lengthy external negotiations he was involved in with environmental regulators and other members of the industry. They sought to set a voluntary goal for emissions reduction that was agreeable to all, but it had important implications for global competitiveness and other environmental and technical ramifications. Tech's insistence that they be given a single and simple target to work to in the near term reflected the fact that, according to an EnviroTech manager, 'they don't believe we have these dilemmas and they don't believe we don't have the power to solve them [internally].'

Members of EnviroTech often recognized the tensions between their work to adapt their practices in order to integrate with Tech and their work in the larger context of environmental considerations. One EnviroTech manager observed that his work involved 'trying to balance a technical solution with a political problem,' recognizing at least some incompatibility between the two knowledge regimes in which he acted.

Conventional level

Finally, at the level of broader conventions (not necessarily bound by the industry), the environmental knowledge regime reflected both the complexity and uncertainty of engagement with the physical ecosystem, and the strong pull to align with the conventions of scalable progress given by Moore's Law. One convention established for governing decisions on the environment is the well-known 'precautionary principle' which states that 'where there are threats of serious or irreversible environmental damage; a lack of full scientific certainty should not be used as a reason for postponing measures to prevent environmental degradation' (Williams, 2001). In other words, organizations should act on incomplete information if the available evidence suggests serious damage to the environment *may* result. This kind of principle, while recognized as desirable by environmental specialists, was an uneasy fit with the work practices at Chipco where, as one Tech manager noted, decisions aren't based on 'a conjecture about the future.' 'If you say there's an X% chance of something happening, the [managers] would say 'come back when you know for sure',' he added.

Other general principles have been articulated through organizations promoting programs like the 'Natural Step' which aims to help companies evaluate their environmental impact by assessing how they stand on four simple 'system conditions' – for example, substances produced by society must not systematically increase in nature, and substances from the earth's crust must not systematically increase in nature. One Chipco environmental manager commented that he felt that the Natural Step was based on reasonable principles, but they were too simplistic and not 'data driven' enough for Chipco. 'We all know the world is going to end', he said, 'the question is when'.

Our review of the ESH (Environment, Safety, and Health) portions of the industry Roadmaps suggests that environmental work at Chipco was also strongly shaped by a pull to conform with the overall conventions of the industry, that is, to establish for 15 years into the future specific and measurable trends for the development of knowledge and technologies. Between 1994 and 2003, the format and metrics of the ESH portion of the Roadmaps shifted considerably, much more so than for lithography or other core process areas over the same years. From the earliest industry Roadmaps the lithography tables have been fully populated with numeric target and the categories for these targets have remained remarkably stable. In contrast, the earliest ESH trend tables were largely text-based and did not contain numeric targets. There is a noted increase in the number of entries with numeric targets over time but there remain many others for which numeric targets are not set. Indeed, of the 11 technology trend tables that do *not* contain predominantly numbers (out of 124 total tables) in the entire 2003 Roadmap, nine of these are in the ESH section. Pressures to display scalable, predictable metrics and deliver on these over time are evident in the changes in the ESH trend tables, but it remains clear that ESH does not behave like other aspects of semiconductor manufacturing technology.

Incompatibility and relative power of the knowledge regimes

Empirically, the semiconductor manufacturing knowledge regime seems to exert a strong pull over the environmental knowledge regime at the work practice and industry level. The first regime is characterized by scaling and divisibility and the second by uncertainty and complexity. But what do these differences tell us? To address this we need to develop criteria that help us talk concretely about these differences as a source of incompatibility, and to speculate on why one regime appears more powerful than the other.

To do this, we have identified the three following characteristics of knowledge in a given regime that enable comparison: the causal specificity of the knowledge, the temporal cycle needed to develop knowledge and the representation of the knowledge through clear and shareable measurement artifacts. By causal specificity, we mean that there is a clear and testable relationship that connects cause to effect in the development of new knowledge. In semiconductor manufacturing, we see this manifested in both scalability and divisibility. While the manufacturing process is immensely complex and the innovation required to continue to shrink component sizes is formidable, the underlying causal relationships are stable, as reflected in the technology Roadmaps. Within Chipco, the continued experience of mastery over the technical challenges leads to a shared optimism about scaling and divisibility. Commenting on ongoing concerns over physical limits to scalability, one manager observed, ‘at every turn in my career people have said there’s a hard limit here and it hasn’t turned out to be true’.

In contrast, the environmental knowledge regime enjoys relatively little causal specificity. Effects – in terms of ecosystem damage, human health, atmospheric change, etc. – have been incompletely understood for years even after causes are identified. Even when a chemical is known to be hazardous and its emission is tightly controlled, as in an example from Chipco when the release of a certain chemical into wastewater was restricted, uncertainty and disputes can surround the connection of cause and effect. In this case, the chemical was tightly bound to another, meaning it was undetectable. Yet would it remain bound? And under what conditions of water temperature, pH, etc? Even if all of these scientific and technical questions are answered, the connection between cause and effect on environmental issues is almost always socially mediated by government regulation or public opinion. Recognizing this, one EnviroTech manager reflected on issues that arose on a project by saying ‘it’s a force – whether it’s real or not – it still has to be worked.’ Of course, establishing causal specificity requires that *others* recognize issues, their causes, and effects as real and this was difficult for EnviroTech to do because of Tech’s high standards for causal specificity.

Second, these two knowledge regimes operated on very different time horizons or temporal cycles for the development of knowledge. The causal specificity within semiconductor manufacturing enabled the creation of an internal pacing mechanism that is virtually unparalleled in any large industry. Building on the predictions of Moore’s Law, the industry has unveiled new manufacturing process generations every 2 years and even been able to develop a mechanism (Roadmaps) for planning new technological milestones with a rolling 15-year time horizon. Of course, significant investment and the shared expectation that scalability has not been exhausted are key factors that link causal specificity to the development of knowledge over short time cycles in this case.

There is no such relentless pacing mechanism within an environmental knowledge regime. Environmental issues come up on their own time, driven by scientific discovery and a slow and unpredictable accumulation of facts. Given the lack of causal specificity as well as ambiguity in terms of where critical information will come from (e.g., data on chemical concentrations in the water under various conditions) it can take long cycles to develop and test knowledge in this regime. Even when the environmental effects of a chemical or material are known, the coordination of actors to address the issues can be difficult and lengthy. One Chipco project was intended to recycle a greenhouse gas so it could continue to be used in manufacturing. However, technical problems during development meant that the recycling equipment would be seriously delayed and would miss its integration with the manufacturing process generation it was designed for. An engineer involved in the project pondered, ‘how long will it take the government to get regulation on [the greenhouse gases], ten years at least, and by then [the manufacturing process] will be using [a new chemical] and it won’t matter’. This illustrates how the weak causal specificity and a much longer developmental time horizon of the environmental knowledge regime put it in a much weaker position in relation to the semiconductor manufacturing knowledge regime because the knowledge produced is less predictable and accumulates more slowly.

Finally, just as they internally control the pacing of their activities, the actors in the semiconductor manufacturing knowledge regime by and large have unambiguous and uncontested artifacts which are used to ‘share and assess knowledge’ (Carlile, 2004) and mark progress by a large set of actors. Again the material aspects of the manufacturing process (scalability and divisibility) make for a set of artifacts that are causally specific and developed on a short time horizon, but more importantly such artifacts are legitimate and easy to share with the large set of actors in the industry. Given the long time horizons to develop more causally specific knowledge in environmental work such clear artifacts cannot as easily be developed or once developed may not be easily shared with a large set of actors. The variability across versions of the industry Roadmaps is evidence that even within the industry consensus is still hard to reach. Further, outsiders to chip manufacturing firms such as government regulators, activists and local communities have very different and often competing measurements that can destabilize the artifacts currently used.

What we now clearly see is the relationship across these three characteristics that not only structures a given knowledge regime, but also explains the relative power of one over another. Causally specific knowledge that is developed and verified over a short time horizon both requires and generates measurement artifacts that are predictable and easily shared. Such factors enable the relatively linear accumulation of knowledge and the generation of shared expectations, leading to a possibly powerful knowledge regime. Knowledge that is less causally specific and developed over a longer time frame produces artifacts that are by comparison less powerful. Here, it is important to consider the wider social context in which the knowledge regimes operate, for one can imagine knowledge that is highly causally specific, developed over a short time cycle, easily shared, and hence accumulating rapidly, but that gets 'trumped' by other considerations like social values. For example, knowledge in some areas of biotechnology has these features, and its rapid accumulation could be severely curtailed if ethical concerns result in legal or other limits on the use of particular techniques.

In the case of EnviroTech, since there are no larger social forces with clear and sharable measurement artifacts that call into question the progression of semiconductor manufacturing technology, it is no wonder that the EnviroTech group adopted the artifacts of the Tech group. These artifacts worked in the short term, but in the end generated negative consequences for environmental concerns and interests. Identifying the gap between the two knowledge regimes, the amount of incompatibility, helps us understand that EnviroTech's adoption of Tech's artifacts did not function as adequate boundary objects (Carlile, 2002). Without an adequate common knowledge, grounded in this case in divergent material realities, the two groups could not fully recognize or address their real differences in language, meaning and interests (Carlile, 2004). And given the amount of incompatibility in this context, a question remains if there exist artifacts or practices that could function at the boundary to help address the gap between these two regimes in the near term; and if not, what would be the requirements and cost of developing ones that could in the long-term?

Discussion and conclusion

By calling attention to the incompatibility between knowledge regimes, we have demonstrated that the true costs of integration across organizational boundaries may be hidden, and that successful integration of knowledge in the short-term can mask costs in the long-term. As suggested above such problems result from a lack of measurement artifacts functioning as adequate boundary objects for actors from different work practices to use as they share and assess each other's knowledge (Carlile, 2004). However, this paper suggests that the ineffectiveness of artifacts as boundary objects can arise from differences that are both deeper and broader than those apparent at the level of work practices. This more comprehensive view adds to the practice based view of knowledge in that it offers an understanding of how nested aspects of the material world, work practice, and broader institutional and industry forces create powerful knowledge regimes. It importantly focuses on the material world as a starting place to begin understanding the structuring of knowledge. Ironically, some who take a practice view of knowledge and knowing only emphasize the social aspects of practice and so can lose track of the material world that provides the originating contours of the practice under examination. Our view provides a 'turtles on up' account that is often lacking in studies of knowledge in organizations, and brings together the material, social and pragmatic considerations that shape knowledge and knowing in practice.

Our focus on materiality acts as a reminder that prioritizing language in many studies of information systems implementation (see, for example, McGrath, 2002) or of organizational and institutional change (see, for example, Creed, Scully, & Austin, 2002; Maguire, Hardy, & Lawrence, 2004) can led to a detachment from embodied practice and the material interests that lay at the heart of many of the debates that are of consequence to these researchers (Dreyfus & Rabinow, 1993). Indeed the majority of this research calls attention to differences in language and framing (Benford & Snow, 2000), but does not begin to identify structural sources of those differences, how big those differences are and, further, what practical steps could be taken to begin filling in the gap that exist in between parties. By calling attention to the material, work practice, and collective aspects of knowledge in shaping the power and persistence of certain knowledge regimes, this research could be similarly applied to bring a greater degree of concreteness to research that either focuses on language or broad institutional processes. Here, the characteristics of causal specificity, time horizon and artifacts allow one to describe the relative differences in knowledge wielded by each group as well as what is missing that could address those differences.

More fundamental incompatibility may be encountered when the conventions used to measure success or progress are vastly different, as, for example, when religious and scientific perspectives are debated. Our analysis likely applies best to domains that are somewhat similar to the one we studied, namely, domains in which integration of technical knowledge is desired. Indeed, when more fundamental incompatibility is encountered, the power of the knowledge regimes need not rest with those that enjoy an accumulation of causal knowledge. As suggested earlier with the biotechnology, social values may even trump regimes that generate causally specific knowledge over short time cycles and have easily shared measurement artifacts.

Our focus on materiality as a potentially important starting point for the structuring of technological knowledge regimes acts as a reminder that material properties produce particular affordances whose consequences may be masked by focusing primarily on work practices, language, or conventions. However, it is equally important to remember that materiality is not the only factor that shapes knowledge regimes. Indeed, without the significant and continued investment in innovation, time and particularly money to overcome the technological challenges of scalability, the semiconductor manufacturing industry would never have exploited the material properties of silicon as relentlessly as it has. The industry conventions and shared expectations then have a recursive effect that reproduces the knowledge regime over time by enabling the continued exploitation of the material properties. Without the material properties, however, no amount of shared expectations and effort could produce such an accumulation of knowledge using the same technological platform and techniques.

A practical implication of developing a more comprehensive understanding of the sources that structure knowledge differently is an ability to improve the design information systems as boundary objects and at the same time recognize the potential limits of any design effort. Developing information system artifacts that function as boundary objects in settings with vastly different knowledge regimes must at least start with adding artifacts that lie somewhere between regimes in terms of causal specificity and temporal cycles for developing knowledge. Of course expecting one group to move toward less causally specific measurements without resistance is naïve and so any effort must consider representing and measuring dependencies that would allow the interests of the two groups to overlap to some degree. In the case of Tech and EnviroTech representing current regulatory constraints as well as potential future regulatory constraints would provide a consequential middle ground between these two groups. In product development settings the use of information systems such as CAD/CAM starting in the 1980s helped provide measures of 'design for manufacturability' at a sufficient level of causal specificity that allowed the more powerful design group to value changes in design proposed by a manufacturing engineer. Without a CAD/CAM system it would have taken weeks or months to develop an adequate understanding of just how the design would compromise the cost of manufacturing and the quality of its output. CAD/CAM simulations made it easier to represent and measure the consequences of the dependencies (i.e. manufacturability, quality) between these two engineering groups that allowed them to transform their respective knowledge and interests to create a better joint outcome (see Carlile, 2004).

Describing which dependencies need to be represented at the boundary between divergent knowledge regimes lies outside the scope of this paper, but these issues have been addressed both specifically and more generally in other organizational contexts (Carlile, 2002, 2004; Østerlund & Carlile, 2005). The most important point to remember in terms of good information system design is that the constraints, the sources that structure knowledge differently, must be well understood. Too often the design of an information system is left to an outside vendor or, more problematically, is shaped by the measurement artifacts of the dominant group(s) inside the organization. Designing a better information system is about the creation of measurement artifacts that sit at the boundary between regimes – at the differences in the causal specificity and the temporal cycle of the development of knowledge – which, like the CAD/CAM example involves engaging groups with historically different amounts of power.

Finally, a unique implication of the semiconductor manufacturing knowledge regime for information systems is that advances in the power and speed of chips have enabled considerable advances in the capabilities of information technologies. Moore's Law not only reflects the material properties of silicon, but also captures the expectations and accumulation of knowledge in the entire electronics and information technology industries. When these expectations about the accumulation of knowledge and their related measurement artifacts come up against other forms of knowledge grounded ultimately in other material realities, such as understanding disease and illness, or even producing consumer goods, powerful opportunities exist. Less recognized are the equally consequential gaps that can remain when the knowledge embedded in these other regimes is incompletely represented by the information technologies that are developed.

A clear limitation of our empirical exploration of knowledge regimes is that it originates from a single case study, although it compensates for this to some degree by examining and comparing more than one level of analysis and source of data. While the semiconductor industry and its scalability and divisibility is clearly not representative of many organizational settings it does provide an 'extreme case' (Eisenhardt, 1989; Pettigrew, 1990) upon which to understand how knowledge regimes are constructed as well as the sources of incompatibility that shape day to day practice and constrain integration efforts. In some settings, for example between different firms or between specialized knowledge domains that have few consequences for each other, the desire for integration will be lower and hence differences between knowledge regimes may be less problematic.

Our examination of the introduction of environmental considerations into semiconductor manufacturing suggests that knowledge regimes are differently structured, and may be incompatible at different levels even if their work practices can be aligned. It also serves as a warning that as businesses try increasingly to address social issues (i.e. environmental, education, community development, etc.) the causal specificity of knowledge, the time horizon required for the development of knowledge and relative power of artifacts used to measure these public goods are often weaker than technical or business knowledge regimes and their respective measurement artifacts. Nevertheless, understanding what structures knowledge regimes differently is an important first step in addressing these challenging problems. By calling attention to the structuring of knowledge regimes across multiple levels – material, work practice, and broader social conventions – we have tried to push the conversation more concretely to identify steps that can begin to address the gaps that such incompatibility naturally generates.

Acknowledgements

We thank Davide Nicolini, Maxine Robertson, Harry Scarbrough, Jacky Swan and other participants in the University of Warwick Business School IKON seminar for their helpful comments on an earlier draft of this paper. We also appreciate the guidance of special issue editors Sue Newell and Bob Galliers and three anonymous reviewers in developing this paper.

About the authors

Jennifer A. Howard-Grenville is an assistant professor of organizational behavior at the Boston University School of Management. She studies how cultural and institutional processes constrain or advance organizational change, with a focus on changes in corporate environmental practice. She is particularly interested in the microprocesses of organizational and institutional change, including questions of knowledge integration and transformation. Jennifer's work has been published in *Organization Science*, *Academy of Management Executive*, *Organization & Environment*, *Law & Social Inquiry*, *California Management Review*, and several edited volumes. She is the co-author of one book, *Greening the Industrial Facility*, with Thomas Graedel. Jennifer received her Ph.D. in Technology, Management, and Policy at MIT, her MA at Oxford University, and her B.Sc. at Queen's University, Canada.

Paul R. Carlile is an associate professor of management and information systems at the Boston University School of Management. Paul focuses on how knowledge is structured differently by localized practices and given that examines what can be done to effectively manage those differences in knowledge across specialized groups who are dependent on each other. This has led his work to focus on the design of boundary objects as representational artifacts used at the boundaries between different groups. Paul's work has been published in such journals as *Administrative Science Quarterly, Information and Society, Management Science and Organization Science*, as well as a number of edited volumes. Paul received his Ph.D. in Organization Behavior from the University of Michigan, his MA and BA at Brigham Young University in the U.S.

References

ALLEN T (1977) *Managing the Flow of Technology*. MIT Press, Cambridge, MA.

BECHKY BA (2003) Shared meanings across occupational communities: The transformation of understanding on a production floor. *Organization Science* **14**, 312–330.

BENFORD RD and SNOW DA (2000) Framing processes and social movements: An overview and assessment. *Annual Review of Sociology* **26**, 611–639.

BOLTANSKI L and THÉVENOT L (1999) The sociology of critical capacity. *European Journal of Social Theory* **2**, 369–377.

BOURDIEU P (1977) *Outline of a Theory of Practice*. Cambridge University Press, Cambridge, U.K.

BROWN JS and DUGUID P (2001) Knowledge and organization: A social-practice perspective. *Organization Science* **12**, 198–213.

CARLILE PR (2002) A pragmatic view of knowledge and boundaries: Boundary objects in new product development. *Organization Science* **13**, 442–455.

CARLILE PR (2004) Transferring, translating, and transforming: An integrative framework for managing knowledge across boundaries. *Organization Science* **15**, 555–568rid.

CREED W, SCULLY M and AUSTIN J (2002) Clothes make the person?: The tailoring of legitimating accounts and the social construction of identity. *Organization Science* **13**, 475–496.

DOUGHERTY D (1992) Interpretive barriers to successful product innovation in large firms. *Organization Science* **3**, 179–202.

DOUGHERTY D (2001) Reimagining the differentiation and integration of work for sustained product innovation. *Organization Science* **12**, 612–631.

DREYFUS H and RABINOW P (1993) Can there be a science of existential structure and social meaning? In *Bourdieu: Critical perspectives* (CALHOUN C, LIPUMA E and POSTONE M, Eds), University of Chicago Press, Chicago.

EISENHARDT KM (1989) Building theories from case study research. *Academy of Management Review* **14**, 532–550.

FOUCAULT M (1980) *Power/Knowledge: Selected Interviews and Other Writings 1972–1977*. Harvester Press, Brighton, Sussex.

GHERARADI S and NICOLINI D (2002) Learning in a constellation of interconnected practices: canon or dissonance. *Journal of Management Studies* **39**, 419–436.

GLASER BG and STRAUSS AL (1967) *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Aldine de Gruyter, New York, NY.

HANSEN MT (1999) The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Administrative Science Quarterly* **44**, 82–111.

HARGADON A and SUTTON R (1997) Technology brokering and innovation in a product development firm. *Administrative Science Quarterly* **42**, 716–749.

HOWARD-GRENVILLE J (2005) The persistence of flexible organizational routines: The role of agency and organizational context. *Organization Science* **16**, 618–636.

Howard-Grenville J (2006) Inside the 'black box': How organizational culture and subcultures inform interpretations and actions on environmental issues. *Organization & Environment* **19,** 46–73.

Huang J, Newell S and Pan S (2001) The process of global knowledge integration: A case study of a multinational investment bank's Y2K program. *European Journal of Information Systems* **10,** 161–174.

Huang J, Newell S, Galliers R and Pan S (2003) Dangerous liaisons? Component-based development and organizational subcultures. *IEEE Transactions on Engineering Management* **50,** 89–99.

Karsten H, Lyytinen K, Hurskainen M and Koskelainen T (2001) Crossing boundaries and conscripting participation: Representing and integrating knowledge in a paper machinery project. *European Journal of Information Systems* **10,** 89–98.

Leonard-Barton D (1992) Core capabilities and core rigidities: A paradox in managing new product development. *Strategic Management Journal* **13,** 111–125.

Maguire S, Hardy C and Lawrence T (2004) Institutional entrepreneurship in emerging fields: HIV/AIDS treatment advocacy in Canada. *Academy of Management Journal* **47,** 657–680.

McGrath K (2002) The Golden Circle: A way of arguing and acting about technology in the London Ambulance Service. *European Journal of Information Systems* **11,** 251–266.

Merali Y (2002) The role of boundaries in knowledge processes. *European Journal of Information Systems* **11,** 47–60.

Miles MB and Huberman AM (1994) *Qualitative Data Analysis: An Expanded Sourcebook* (2nd edn). Sage, Thousand Oaks, CA.

Murphy C, Kenig G, Allen D, Laurent J and Dyer D (2003) Development of parametric material, energy, and emission inventories for wafer fabrication in the semiconductor industry. *Environmental Science & Technology* **37,** 5373–5382.

Nelsen BJ and Barley SR (1997) For love or money: Commodification and the construction of an occupational mandate. *Administrative Science Quarterly* **42,** 619–653.

Nonaka I (1994) A dynamic theory of organizational knowledge creation. *Organization Science* **5,** 14–35.

Orlikowski WJ (2002) Knowing in practice: Enacting a collective capability in distributed organizing. *Organization Science* **13,** 249–273.

Østerlund C and Carlile P (2005) Relations in practice: Sorting through theories on knowledge sharing in complex organizations. *The Information Society* **21,** 91–107.

Pettigrew AM (1990) Longitudinal field research on change: Theory and practice. *Organization Science* **1,** 267–292.

Scarbrough H, Swan J, Laurent S, Bresnen M, Edelman L and Newell S (2004) Project-based learning and the role of learning boundaries. *Organization Studies* **25(9),** 1579–1600.

Shell J and Scholz R (2003) Consensus building and knowledge integration methods applied to technology implementation in the Swiss waste incineration industry. *International Journal of Environmental Technology and Management* **3(3/4),** 368–389.

SIA (1994) *The National Technology Roadmap for Semiconductors*. Semiconductor Industry Association.

SIA (2001) *International Technology Roadmap for Semiconductors*. Semiconductor Industry Association, http://public.itrs.net/Files/2001ITRS/Home.htm.

SIA (2003) *International Technology Roadmap for Semiconductors*. Semiconductor Industry Association, http://public.itrs.net/Files/2003ITRS/Home2003.htm.

Star SL (1989) The structure of ill-structured solutions: Boundary objects and heterogeneous distributed problem solving. In *Readings in Distributed Artificial Intelligence* (Huhns M and Gasser L, Eds), Morgan Kaufman: Menlo Park, CA.

Suchman L (1987) *Plans and Situated Actions: The Problem of Human Machine Communication*. Cambridge University Press, Cambridge, U.K.

Thévenot L (2001) Pragmatic regimes governing the engagement with the world. In *The Practice Turn in Contemporary Theory* (Schatzki T, Knorr centina K, and Savigny E, Eds), Routledge, New York.

Tyre M and von Hippel E (1997) The situated nature of adaptive learning in organizations. *Organization Science* **8,** 71–84.

Van Zant P (1997) *Microchip Fabrication: A Practical Guide to Semiconductor Processing* (3rd Edn). McGraw-Hill, New York.

Williams J (2001) Australia State of the Environment Report. http://www.deh.gov.au/soe/2001/biodiversity/glossary.html.

Zollo M and Winter S (2002) Deliberate learning and the evolution of dynamic capabilities. *Organization Science* **13,** 339–351.