Hoyles C. & Noss R. (2007). The Meanings of Statistical Variation in the Context of Work. in Lesh, R., Hamilton, E. & Kaput, J. (Eds.), *Foundations for the Future in Mathematics Education*. Mahwah, NJ: Lawrence Erlbaum Associates. (pages 7-35)

THE MEANINGS OF STATISTICAL VARIATION IN THE CONTEXT OF WORK

Celia Hoyles & Richard Noss

Institute of Education, University of London

Abstract

The changing nature of working practices and the ubiquity of computationally-based systems has brought with it the need for many employees to develop understandings of the IT-based models that are increasingly part of their working practice. Our goal in this paper is to sketch some of the categories of mathematical knowledge that characterize working practices in an era of ubiquitous technology, and to explore the meanings of this mathematical knowledge – specifically statistical variation – at the boundaries between the different communities of actors within them.

Introduction

In his stark analysis of modern working practices, Reich (1991) classifies working practices into three types; symbolic analysts, who solve, identify and broker problems by manipulating symbols, *in-person* services (jobs based on interaction with people) and *routine* producers. His thesis is that only a relatively few workers will contribute materially to the knowledge economy, and he paints a picture of the kind of education that is appropriate for such workers:

> Symbolic analysts solve, identify and broker problems by manipulating symbols. They simplify reality into abstract images that can be rearranged, juggled, experimented with, communicated to other specialists, and then, eventually, transformed back into reality (Reich, 1991).

We leave aside Reich's broader economic concerns here, as they fall squarely outside of our expertise. Reich's thesis is, however, broadly convergent with other analyses (see, for example, Zuboff, 1988). The implications of this thesis for the kinds of knowledge required by operatives within such workplaces are relatively clear-cut. A classic case is reported by Raizen (1994) who cites Lesgold's work with USAF technicians: the more highly skilled mechanics differed from their less skilled colleagues, chiefly in their ability to hold better "mental models of the systems they were working with (how components work, what their functions are, and how they relate to the system as a whole)…" (p. 73).

The key point about symbolic analysts is that they model the processes of work into abstractions that can be communicated to other specialists. Whether or not Reich is right in the detail, his central contention that it is the manipulation of symbols rather than things is beyond doubt. The question, however, is not only to understand what the symbol-analysers need to know to develop useful symbolic representations of the workplace, but what the implications are for those who 'consume' the results of the symbol-analysers. In particular, in the years that have elapsed since Reich proposed his analysis, the relationships and intersections between different layers of the workforce have shifted in significant ways. Part of our agenda in this chapter is to seek to reshape our understandings of who needs to know what, and the ways in which the knowledge takes on a different character in relation to the sub-communities of practice within workplaces.

The invisibility of knowledge at work

First we consider a major epistemological and methodological difficulty, which concerns the *invisibility* of knowledge at work. From a socio-cultural point of view, the invisibility of knowledge required in modern work has been commented on by many authors (see, for example, Nardi and Engeström, 1999). From this perspective, invisibility is seen to apply to entire swathes of working practices, such as cleaning or maintenance, rendering the individuals and communities concerned unnoticed and ignored, their place in the chain of work-related activity disregarded. For example, Bishop (1999) argues that distinguishing between visible and invisible work affords an understanding of the ways employment relations are emerging in post-industrial societies, and calls into question a set of implicit assumptions that all invisible work 'should be made visible'. From a rather different orientation, Star and Strauss (1999) consider instances that exhibit a range of indicators of what counts as work and how it appears in different situations. These include the creation of a 'non-person' (in which the employee but not the work is invisible), taken-for-granted work that they characterise as 'disembedded background work' in which the workers are visible, but their work is not, and the abstraction and manipulation of indicators – a category that comes closer to our own concerns, in which both work and people come to be defined as invisible. This last class emerges in two cases:

> i) Formal and quantitative indicators of work are abstracted away from the work settings, and become the basis for resource allocation and decision-making.
>
> When productivity is quantified through a series of indirect indicators, for

example, the legitimacy of work may rest with the manipulation of those indicators by those who never see the work situation first hand.

ii) The products of work are commodities purchased at a distance from the setting of work. Both the work and the workers are invisible to the consumers, who nonetheless passively contribute to their silencing and continuing invisibility (Star & Strauss, 1999).

While questions of legitimacy and silencing are important and challenging facets of the invisibility of work, our concerns are primarily with the visibility of *knowledge* held by persons or communities (rather than of persons themselves or the products of their labour). Knowledge, of course, does not come out of nowhere: it is embedded in people, in contexts, and in situations. Nevertheless, there are, we believe, important aspects of the problem that cannot be captured without a corresponding focus on the ways that knowledge in general and mathematical knowledge in particular – not just people – has been transformed in the modern workplace.

Concerns with the ways that mathematical knowledge is used in workplaces predates the demographic and organisational shifts engendered by technology and has tended to pay more attention to the visibility or invisibility of the knowledge, in contrast to research in the socio-cultural tradition, as mentioned above (for an overview of research in mainly non-technical settings, see Bessot & Ridgeway, 2000). Mathematical knowledge is judged as invisible as it tends to be deeply embedded within the representational infrastructures of the models, tools and artefacts of the workplace. Our own work with bank employees, nurses and engineers has evidenced this invisibility and has shown how mathematical knowledge in use is characterised by fragmented and pragmatic strategies intertwined with meanings of the mathematical knowledge and situational “noise”; and that mathematical knowledge is transformed when it crosses boundaries between different situations or settings (Hoyles & Noss, 1996; Hoyles *et al*, 2001; Noss *et al.* 2002; Kent & Noss, 2002; for a summary see Noss, 2002). Others have reported similar outcomes following research in technical workplaces, such as the automotive industry, Smith & Douglas, 1997; structural engineering, Hall & Stevens, 1995 and Hall, 1999; and industrial chemistry laboratories, Wake & Williams, 2001 and Williams & Wake, 2002). It must also be noted that the invisibility of mathematical knowledge is compounded by the general perception that mathematical and technological competencies consist of sets of decontextualized skills or techniques, disconnected from each other, and from their context of application (see, for example, Smith & Douglas, 1997; Clayton, 1999), an issue that may be generational differences in interpretation (see Zevenbergen, 2004).

All these studies have had to face the methodological challenge of making visible the embedded mathematics of the practice in order to study and analyse it. Most have undertaken ethnographies, often focussing on "disruptions" in the routines of work or on communication across different representational infrastructures (Hall, 1999; Noss *et al.* 2002). Yet despite its invisibility, Hall et al. (2002) have argued that what makes subject-matter knowledge systems powerful is also what makes them difficult to study; it is powerful knowledge that is both widely distributed and massively influential in shaping the ways people interpret their working activities, while at the same time, profoundly embedded in the representational infrastructures that permeate working practices. From a methodological point of view, Hall et al. – like us – regard enhancing the visibility of subject-matter knowledge as a crucial analytical challenge.

We therefore now provide a brief overview of the methods used in our previous and ongoing studies, which we use to provide the illustrative empirical data in this chapter.

Methods

In each new study of a particular work sector or factory, we have established a modus operandi that consists of five main components: interviews (which may be telephone, audio-recorded) with site or technical managers, site visits (involving work shadowing, and impromptu interviews where appropriate), iterative analyses (ongoing throughout the site visits), cross-factor analysis to draw out common components, and validation of analyses (with stakeholders within the work sector). Initially, a list of companies in a particular sector of work is drawn up following consultation with and advice from the relevant professional organisation, from which a sample for case study is selected by the project team to represent as far as possible the spread within the sector and geographical locations. It is worth noting that there are inevitably some problems in obtaining and retaining access to companies. Each case study then comprises an initial semi-structured telephone interview with a key person in the company, following a set of agreed questions. The aim of the interview is to obtain a broad picture of the work and the range and profile of the employees in terms of skills and expertise, but most crucially to try to identify one or more people for observation and interview on subsequent site visits. Such people might be the Technical Manager or somebody who is in charge of several teams of people. Once identified, data from the telephone interview provide information on the job title of the identified person, what work is involved and where he/she fits into the company structure.

Following the interview, site visits are arranged. These are similarly structured according to sets of pre-defined questions, but with the flexibility to follow up interesting avenues for further investigation, including interviews where possible. This is particularly important as companies give different titles to jobs that might appear the same to the outsider (team leader, shift manager, group leader, process technician), and describe what we might see as mathematical processes in a different language from that of mathematical discourse. After general orientation in the company, we seek to tease out through questions and observations (some gleaned from work-shadowing) how mathematical knowledge might play a role in the workplace and specifically, how artefacts are used and how they might mediate the knowledge required. If possible, we try to document the extent to which employees are expected to display mathematical problem solving skills, for example, appreciate links between disparate data sources or representations, base decisions on underlying structures, models or logical analyses alongside immediate concerns. We similarly probe for specific examples of when a 'skill' might be needed. In this endeavour we often use critical incident techniques and ask for a description in as much detail as possible of an occasion when something critical happened, something did not work out, or there was a dispute as to appropriate action, how and why it arose, and how it was resolved. We paid particular attention to the roles of artefacts, especially computational objects and systems.

Case studies are then drawn up, and provisional versions sent to each company to check for accuracy and to safeguard confidentiality. The next step is to draw out common themes or issues that span case studies from one sector, and could be seen as ‘characteristic’ of a typical type of work in the sector. Finally, we attempt to involve professional associations in organising presentations of research findings to appropriate employers and sector representatives allowing for wider validation and if necessary, modification. We are sensitive to the fact that the case studies do not represent a random sample, either of a sector or even of a type of business within a sector. These validation exercises help us see limits to the generalisation of our findings as well as assisting us in developing a language to communicate to the relevant work communities.

We are well aware that asking managers to talk about the knowledge requirements of their employees' work is not at all the same as asking the employees, observing the employees at work, or studying longitudinally the activity system and the employees' actions within it. Nonetheless, this is always the first step of our studies and we will present many extracts from interviews with managers. We ask the reader to bear in mind our objective: it is to elaborate our theoretical framework, derived from a substantial corpus of data emanating from several studies in a variety of workplaces – and we use these extracts as illustrations of the theory as it materializes. Viewed in this light, the way that managers describe their knowledge requirements for example does, we think, provide an illustrative backdrop for the elaboration of our emergent theoretical position, by contrasting, for example, their perspective with that of other workers.

Information technology and the modern workplace

In many manufacturing companies in modern workplaces, huge amounts of data are gathered daily to measure different aspects of the work process. For example, in the manufacturing sectors considered in the Mathematics Skills in the Workplace project (Hoyles, Wolf, Kent & Molyneux-Hodgeson, 2002), it was reported that the companies surveyed were hugely concerned with improving efficiency and production, and data were gathered and manipulated to provide (potentially) an evidence base for action planning and decision making. Although company practices varied (with the larger companies proceeding more 'scientifically'), it was found that a common feature across sectors was the expectation that most employees, regardless of level, were to be involved in the process of increasing productivity or producing efficiency gains. Therefore the study concluded that all (or nearly all) employees needed relevant and interrelated mathematical and IT competencies.

With increasing emphasis on process control and improvement, and the need to deal with more complex contractual arrangements with customers, data – often represented on a spreadsheet (see as an example, Figure 1) – were used to monitor and analyse a factory's performance at different levels (e.g. on single production lines, or divisions of the factory) and over different timescales (ranging from dealing with day-to-day technical problems, to long-term planning of production improvements up to several years ahead). Two points are worth making at the outset. First, the columns of the spreadsheet represent the variables of the process, some of which are "Key Performance Indicators", KPIs, such as downtime for machine (that has to be minimized), or number of products per hour (which, of course, normally has to be maximized). A second point is that these KPIs are seldom independent and many are connected by a specified (but usually opaque) mathematical formula.

**Insert Figure 1 about here **

In the Hoyles *et al*. (2002) study, team leaders were identified as playing a pivotal role in the operations of the factory, "reading in" information to the abstract model (such as that underlying the spreadsheet) for upward communication to management, and "reading out" information from the model for dissemination to operational workers. To fulfil the latter purpose, data related to KPIs were often displayed on factory floors, in the form of spreadsheet-derived graphs, which also depicted output related to given targets (see Figure 2). Team leaders also often had to identify poor performance, and recognise anomalies and erroneous input data.

**Insert Figure 2 about here **

At the same time, managers were increasingly required to take decisions based on output from these quantitative models and to extrapolate trends, while accounting for other, less-easily-quantifiable information, such as the production capacity of the work force and the current "industry climate" (see Hoyles *et al.*, 2002). Apart from the vertical metaphors of upward and downward communication, such information also had to be channelled *horizontally* within a project team to communicate among people with different skills, and *outwards* to customers.

We now proceed to elaborate the new literacies required in the workplace, by presenting some examples using some of the case study extracts that appeared in Hoyles *et al.* (2002). We begin in the company within which the spreadsheet shown earlier was used. It is reasonably typical: a large packaging company employing about 220 persons, that made printed cartons on a contract basis. Contracts could be stand-alone but increasingly there were multinational contracts for corporate clients covering many countries and different types of cartons. Cartons are printed, cut out and creased on very high-speed, computerised printing presses. Typically, carton design is done by the client, but the company has to work out (nowadays, using computer software) how optimally to 'fill' the individual carton designs onto the raw cardboard sheet. Cartons leave the factory as flat sheets, which are then formed into box cartons at the client's packing plants.

First, it was reported that a change had occurred in that key managers of information in many companies were now the team leaders who supervised shop-floor operations, a crucial change instigated by the introduction of IT-based technology. Thus team leaders needed to manage and communicate information much more than in the past. A senior manager told us:

> Manager 1: We're data-driven much of the time, and it is clearly our strategy to push a data-driven approach right down through the organization. Now, team leaders have to come and present to me a lot of analytical data about what happened in their shift. Five years ago it didn't happen …

We asked M1 if he could describe an employee who had found it difficult to cope with this complexity in the light of the introduction of massive computerisation:

> M1: We had somebody, X, who had been in the section a long, long time. He could not manage the six machines, the complexity of it… He found it very difficult to be able to juggle the dynamics of what was happening... He could do the mechanics of it because he knew how to do that... but he was not able to interpret it, to transform it into a trend or a problem solving analysis. He wasn't able to really use the information in a well-constructed argument. He could not present an argument which was fact-based, by using the numbers and the information that we do collect. …not really understanding what that data means or how it impacts on the business.

The second change was in the marketplace that was more competitive, with profit margins very much leaner than they used to be, and clients demanding higher quality and more complex contractual arrangements. As M1 commented:

> M1: If you look back 10 years, the variability in our products would have been huge. And the customers didn't care so much, because their supply chain was not that sophisticated, and waste was not very visible; the drive for profit and cost reduction was nothing like what it is now.

Managers thus needed to determine realistic performance criteria that formed part of complex models of production, including for example, product output as a function of 'man-hours', productive machine time, and relative costs of storage, design and transport. They then needed to use these data to assess performance against criteria and targets. Fluent manipulation of variables, some of which were quite invisible, was therefore crucial. Thus management needed to see data both as a concrete representation of what the company produced, and, at the same time, as an abstract representation that identified the implications of the data in terms of general trends that had both historical and predictive significance. As part of this process, they needed to be able to communicate complex information from interrelated data sets, spot errors and troubleshoot, as illustrated by a second manager we interviewed in another manufacturing company:

> M2: We record everything on a log and it actually tends to come to me as well as everybody else and we keep our own Excel spreadsheets which are followed over time and it is obvious if say you have a bad measurement.
>
> … It just leaps out at you. Because obviously the real problem is that if you've got bad measurements and you don't realize that there is something funny about the numbers then you go to the next stage in the analysis. Well the software doesn't know it is a bad measurement so it just throws it in.

So how might we describe what team leaders and managers must ‘do’, and what, for example, X could not do? It is not helpful to describe this purely in terms of ‘competences’, of say, understanding variables and trends. This does focus on knowledge, but misses a crucial aspect of what it means to use knowledge *in situ*, where any understandings are necessarily articulated within the workplace discourse, and supported by contextual cues within it. The knowledge mobilised in use is not codified knowledge (Eraut, 2004), in the sense that it is readily available in published sources. It is knowledge that is mobilised in the practices of working cultures, and its realisation in activities is achieved in intimate connection with practices and artefacts of the workplace, specifically with the technologies that perform and monitor the work process.

Given the need to characterise the kinds of knowledge involved, we have coined the description *techno-mathematical* knowledge, to highlight the fusion of IT and mathematics. It is obvious that there is huge variation in the depth of this techno-mathematical knowledge required by different elements within and between working communities. From the perspective of Reich's symbol analysts, the kinds of knowledge required are strongly associated with academic knowledge taught, for example, in university departments of mathematics or theoretical physics. It would certainly be worthwhile to explore the extent to which these curricula are actually aligned with the requirements of modern working practices: but this is not our concern here. Instead, we focus on individuals and communities within working practices that are relatively uneducated in these domains, yet who have to work with their products. For this reason, we focus not on techno-mathematical knowledge *per se,* but on techno-mathematical *literacies,* what needs to be appreciated, interpreted, and communicated horizontally and vertically within working communities. The term 'literacies' is appropriate, in that the kind of knowledge required is analogous to that required in relation to reading and writing – not to write books but to read them, not to construct literature but to talk about it, and think about its meanings. *Techno-mathematical Literacies* (TmL) are technically-orientated functional mathematical knowledge, grounded in the context of specific work situations and employed as devices for making sense of the systems that pervade and operationalize them. Part of the research endeavor in our current project, *Techno-*

mathematical Literacies in the Workplace[1] (henceforward referred to as the *TmL Project*) is to characterize the different meanings associated with TmL as articulated by the different communities who use them in a variety of working practices.

Conceptualising techno-mathematical knowledge: boundary objects and situated abstraction

We have argued that the work process has increasingly become mathematized through models instantiated in different technical artifacts, with the outcome that huge amounts of quantitative data have to be interpreted as a basis for action and decision-making. The question then arises as to what happens to individual and collective knowledge under these conditions and how it is mediated by the technology. This in turn requires a more analytical description of the knowledge, who uses it and for what purposes.

We suggest, following on from the ideas of 'disruption' mentioned earlier, that one method to obtain insight into how TmL is used at work, is to consider how the relevant knowledge is perceived across *boundaries* within the workplace, as it is at these boundaries that ideas may be contested, disruptions may occur, and different languages of description have to be aligned. We conjecture that there are at least two major trajectories of what it means to cross boundaries, the first concerns people and the second objects.

The first phenomenon, which we will call *devolution,* describes how many individuals need now to understand aspects of the system that previously would not need to have been understood by an

[1] The research project "Techno-mathematical Literacies in the Workplace" is funded by the UK Economic and Social Research Council as part of the Teaching and Learning Research Programme [www.tlrp.org], Award Number L139-25-0119. Some of the preliminary work reported later in this chapter was conducted in collaboration with Phillip Kent, whose contribution we wish to acknowledge.

individual within that community (it would doubtlessly have had to be understood by *someone* even if they were never physically part of any community within the workplace). This is, perhaps, a surprising spin-off of computerization, given that the initial impetus for the introduction of digital technologies was often unashamedly that of deskilling the workforce both quantitatively and qualitatively. Yet, as one manager in a packaging plant recently commented to us: "We replaced six operatives by two, but then we found out that the two remaining people needed to understand a lot more than before!”.

The second trajectory, *communication*, names the way in which people need to share knowledge with others across boundaries of space and time. This *horizontal* dimension of knowledge – and therefore learning – has been theorized by Engeström (1996), who argues that a major challenge for individuals and social groups is learning how to cross the social and cultural borders between different activity systems (see Guile and Young, 2003). For Engeström, these 'boundary crossings' point to a *horizontal development* that is to be seen as radically different from the traditional *vertical* notion of mastering a skill or a hierarchical body of knowledge. The second aspect is thus about objects that populate the boundaries and facilitate communication across it, and it is these objects that we now discuss.

Objects at the boundary: windows on devolution and communication

A boundary object (Star, 1989; Star & Griesemer, 1989), names an important class of knowledge artefact shared between different communities of practice and can be used differently by the communities.

> Boundary objects are objects that are both plastic enough to adapt to local needs and constraints of the several parties employing them, yet robust enough to

> maintain a common identity across sites… Like the blackboard, a boundary object 'sits in the middle' of a group of actors with divergent viewpoints (Star, 1989).

Thus boundary objects can be shared between different levels of an activity system or between different activity systems. Each community might view the boundary object differently, talk about it in distinct ways, express its functions diversely or see different parts of the system in it, but nonetheless the boundary object can provide a means to think and talk about an idea and negotiate its meaning. Boundary objects might, for example, comprise a symbolic map of the work process thus opening it up for discussion (devolution), or sit at the intersection between team leaders and managers (communication).

Our assumption is that in the modern workplace, many boundary objects include symbolic representations of mathematical relationships and thus make it possible for them to be interpreted mathematically. Although representations create a layer of invisibility with respect to the mathematical knowledge underpinning them, they simultaneously – and contradictorily – afford a common visible framework that allows diverse communities to act and think as if they had a common purpose. We suggest that this common symbolic discourse can be a starting point for negotiation of meanings. For example, the workplaces described earlier were rich in mathematical representations of process, abstractions of the systems that characterised production, portrayed through spreadsheets or graphical charts. If we consider these representations as forming elements of boundary objects, the techno-mathematical knowledge required to reason with them depends on three factors: first on the way they are represented, whether as lists of figures (as in a spreadsheet) or as a dynamically-represented graph; second, on the uses to which they are being put in the workplace; and third on the experience of the individual and the extent to which appropriate mathematical knowledge might be mobilised. For example, if we consider the spreadsheet illustrated earlier, we know that some columns of data are related mathematically to others. This implies that the variables in the dependent column cannot be manipulated independently of their parent. While most employees would not be required to identify the mathematical relationship involved – and certainly would never be required to express it algebraically – they *may well* be required to understand the basic shape of the relationship (is it direct or inverse?), the effect of a change in one variable on others, and the identification of any anomalies on the basis of the underlying relationship - but of course only in so far as these understandings are meaningful in the workplace and not simply as abstract entities.

Individual and collective knowledge at the boundary

From an activity-theoretic perspective, boundary objects facilitate the achievement of an object (goal) in the face of competing interests from different communities (Bowker and Star, 1999), and in contrast to our approach, the knowledge within each community and the way it is recontextualized is not a central concern. Similarly in using activity theory in the context of employee training, knowledge is not explicitly considered: as Guile and Young (2003) put it, "while activity theory stresses how activities are needed that enable students to re-locate what they already know as a step to acquiring new knowledge, it has little to say about the form or content of this knowledge".

For our research agenda, a major interest is precisely the study of the meanings attached to boundary objects, which will include the meanings associated with any mathematical representations. It is worth noting that these representations may not be expressed using the official symbols of mathematics, but rather with *ad hoc* representations of some specific aspect of factory production. Thus we need a language that catches what is simultaneously situated and abstract about the mathematical knowledge involved. We employ the idea of *situated abstraction* to pay equal attention to the notion of knowledge and to the language in which it is expressed: a situated abstraction is the expression of a mathematical abstraction framed by the artefacts and the discourse of a community. Any mathematical models involved may be different from those that are conventionally understood: not necessarily mediated by formal mathematical symbol systems and artefacts, but rather by 'situated' techno-mathematical artefacts (see Hoyles & Noss, 1992; Noss & Hoyles, 1996). Thus we would expect workers using technological tools as part of their practice, tools that mediate the relationships that structure production, to express and communicate their knowledge in ways shaped by these tools.

Situated abstractions are the means by which mathematical ideas can be communicated to self and to others – without necessarily having recourse to conventional mathematical symbolism. Such representations are not to be conceived as an optional extra, an illustrative form that embellishes or gives voice to an already-formed object of idea: on the contrary, we take it as read that artefacts and symbol systems are *constitutive* of meaning and – by implication – of thought (see Sfard, 2001, for an insightful perspective on this question). Thus, the meanings that people hold of mathematical ideas are framed by the representational infrastructure with which they are expressed. So for the spreadsheet shown earlier, the column of numbers stands for a variable in a process which itself stands for a rate, a quantity or even a relationship. Moreover, the meaning of the mathematics is inextricably connected to work process activity. For example, a variable which 'represents' the runtime of a machine, is not to be thought of just as 'time', but understood in relation to the actual machine, its history and its idiosyncrasies.

The traditional view is that workplace knowledge is pragmatic in character, and that there is a natural antagonism between pragmatic and theoretical knowledge, in terms of its purposes and its forms of representation. The notion of situated abstraction calls into question this distinction, suggesting that the seeds of the theoretical are present in the pragmatic, especially in situations where the tools and artefacts involve symbolically represented objects.

In order to give shape to these ideas, we now elaborate our example of TmL using some extracts taken from our ongoing work in the TmL project, derived from case studies of factories involved in mass production and process improvement.

The meanings of variation at the boundaries

A significant number of companies across the sectors studied in Hoyles *et al.* (2002) were beginning to invest in operating methodologies concerned with quality and continuous improvement. These methodologies go beyond 'conventional' efficiency improvements of the type described earlier. They seek through a strict testing schedule based on statistical theory to establish higher absolute levels of quality, with smaller margins of error and tolerances. In the past, any such testing regime was both designed and interpreted by senior managers and engineers/technicians, with team leaders simply being required to collect information and report problems to more senior colleagues. This trend is only a statistical edge of the current we outlined earlier, but it has a very clear connection with a body of explicitly mathematical theory, and we are concerned to understand just how much of this – if any – is relevant, in what form, and for whom. So the specific example of TmL we now consider in more detail is that which is called into operation in companies that are striving to use a statistical process control (SPC) methodology such as "Six Sigma"[2].

In one such company, a pharmaceutical company, we were told in an interview with a manager that the main task was to reduce variability. Using as an example the target weight of a tablet, he explained:

> M3: The whole ethos of manufacturing is to eliminate variability, then you won't have defects and your customers will be happy. When you have poor process control, you have to compensate for it, maybe by increasing the average weight of material in the product (in this case the tablet) [*so the normal distribution lies always sufficiently above the lower specification limit – CH/RN*], but if you can reduce the sigma, you can use less material and make the product more cheaply. So, all our efforts involve asking 'why is the variability like this?', 'how can we close the gap?'

M3 stressed that all layers of the workforce, not just managers, must become SPC-literate and this meant improving understanding of *variation*, since it is this understanding that is thought to play a crucial role in improving practices at work. As Joiner puts it in his book on Fourth Generation Management:

---

[2] Six Sigma seeks to control manufacturing processes so tightly that only three products per million will be defective; that is, in the distribution of any measure of the product, there are three standard deviations (sigmas) either side of the mean value that fall within the quality specification limits.

> Our ability to produce rapid, sustained improvement is tied directly to our ability to understand and interpret variation. Until we know how to react to variation, any actions we take are almost as likely to make things worse, or to have no effect at all, as they are to make things better (Joiner, 1994).

But what does ‘understanding variation’ in the context of work mean? Wild and Pfannkuch (1999) have argued that: “variation is omnipresent; variation can have serious practical consequences; and *statistics* give us a means of understanding a variation-beset world (our emphasis).” But can people without any background in or knowledge of statistical thinking gain understanding of variation from experiences in their workplace insofar as they are required to characterise this variation, quantify it and seek to reduce it?

Konold and Pollatsek (2002) suggest that understanding data means appreciating the existence of a ‘signal in noisy processes’, a stable value, the central tendency of a data set, with (inevitable) variation. They argue that “implicit in our description of central tendency is the idea that even as one speaks of some stable component, one acknowledges the fundamental variability inherent in that process and thus its probabilistic nature”. Because of this, they claim that “the notion of an average understood as a central tendency is inseparable from the notion of spread” and that “average and variability are inseparable concepts”. If, as Konold and Pollatsek suggest, rather little attention is paid to these ideas in school, is it possible that workers can appreciate them, let alone come to use them effectively?

We seek to illuminate these issues by presenting a further series of extracts that start with the notion of a boundary object, and follow with descriptions of the meanings that different communities have expressed about these objects.

The SPC chart as a Boundary object

The first step factories take when moving to process control, is to select key performance indicators, KPIs, of productivity and efficiency, along with a set of standardised in-process control measures to assess them. KPIs are the output of a process, and the data used in SPC are the measures of these KPIs over time. What do workers need to know about data analysis in order to monitor performance? What are the meanings they assign, for example, to the central tendency of a particular KPI process and the variation that might be evident? Fundamentally important in SPC is the drive to control variation and central to achieving this control is to distinguish between common cause variation and signals or 'special cause' variation[3]. Common cause variation is ever present in any process or interconnected set of processes. Special causes arise for reasons *outside* the usual process. They can contribute either small or large amounts to the total variation, but typically have a much larger impact on variation than the common causes – and of course special causes affect the central tendency of the data set:

> The important point about common cause variation is, it is argued, that there is no single cause: it is the result of a set of interacting factors. So, to change common cause variation requires reconfiguring the entire system. Special cause variation, on the other hand, demands immediate investigation with – if possible – an immediate remedy. In order to reduce common cause variation, a much more subtle and in-depth appreciation of the process is required" (Joiner, 1994).

[3] Signal seems to be used in the context of work rather differently from Konold and Pollatsek (2002), here we use the term as we heard it used in the workplace.

If an SPC-literate workforce is needed to implement process improvement, how do companies proceed? First, our case studies indicate that they seek to make the variation in a data set *visible* in the form of SPC charts, which are produced automatically for each KPI. SPC charts have three major features: data are plotted in time order; a centre line – typically the mathematical average or central tendency of the process (as characterised by Konold and Pollatsek, 2002) – is drawn; and finally statistical control limits are added that indicate the width of the common cause variation, the historical expectation of the normal variation. These upper and lower control limits are based on what the process is capable of doing, calculated by statistical formulae used on historical data sets. Depending on the nature of the data, other features are sometimes added, namely a target production figure and the upper and lower limits of a customer specification (the limits of quality that are acceptable for a particular product).

A simple example SPC chart is shown in Figure 3 showing the 'historical mean', the upper and lower control limits, the customer specification and the target line just above the mean. The chart shows one point outside the lower control limit that signals a special cause.

Insert Figure 3 here.

In the factories we visited that used SPC methodology, SPC charts could be seen on the shop floor and in management offices. They are familiar ‘mathematical’ objects, regarded as important and taken as the focus of discussion in regular team meetings. The data are meant to be transparent in that they are agreed measures of process and they are input by operators in the factory – although the chart itself and the control lines on it are produced automatically. An SPC chart can be thought of as a boundary object comprising a knowledge artefact based on a mathematical and statistical representation, which might be interpreted differently depending on one's position in the activity structure or degree of mathematical knowledge. What are the meanings that different groups of workers in the factory have of these charts? How far are they aligned with each other and with statistical interpretations? How are they mediated by the technological tools that create the charts? How do the charts provide insights into the factory process?

We present below some different descriptions of the meanings of the SPC methodology and of the SPC chart as related to us by key people in process improvement: first, a company statistician; second, by company managers; and finally, by team leaders working on the shop floor.

The Company Statistician

The statistician's (S) point of view sets out the rationale of SPC. His view, expressed in no uncertain terms, was that everybody working in the company needs to understand the meanings of variation and uncertainty in the context of SPC, which he argued could be achieved rather simply but rarely was evident:

> S: What I'd like people to do, whether they're a new manager, a new scientist, a new operator or whatever, is to arrive here understanding variation and issues of uncertainty … there's no reason that I can see why everybody shouldn't arrive here with that knowledge. It's very simple knowledge. But hardly anybody does.

S argued that it was crucial that employees distinguished work outcomes along two dimensions: "what we *want* to get?" (the customer specification) and "what we *expect* to get?" (data varying between the two control limits). Taking effective action demanded that workers compared and acted upon actual results along *both* these dimensions with the desired state in any process being that the expected variation was within the variation set by the customer specification. At the same time no action should be taken when the data points were simply the result of random variation.

S went on to argue:

> S: …. There is a great temptation for people to see a number that's lower or higher than the previous one and feel they have to make a change. Deciding this is not really evidence for change therefore I don't need to alter the process. ….Anybody going into any industry should believe that this variation issue is important for them and it's important in all their decision-making.

So the first key production technique was not to act on changes that were simply the outcome of common cause variation, even if this meant the customer specification was not achieved:

> S: …A key part of using these charts is you need to know which strategy for improvement to use. A huge proportion of the world *investigates a point outside the customer specification line (a disappointment) and tries to improve* [italics added]. And that is very wasteful and inefficient.

The second technique was to investigate signals, deviations from expectations, explain them and then either correct problems designing them out of the process or planning to incorporate them into the process. In both cases, the aim must be to stabilize process as soon as possible. Reaction to signals was important but still more so was improving the process that is reducing variation. As S put it:

> S: Variation has costs, and reduction in variation constitutes an improvement. Thus once the process is stable, it is important to start proactive improvement.

To improve, that is to ensure the control limits are within the customer specification *and* variation is reduced, requires understanding common cause variation, analysing the multiple factors that give rise to it and then seeking ways to reduce it. As S explained:

> S: You'd have to get much more involved in fully studying your process and understanding where that variation is coming from. You'd use all of the data points and lots of information, with expert people to facilitate the work. It could be anything to do with the variation in the raw materials. It could be variation in the environment, it could be differences in the machines you use, it could be differences in the way operators work, it could be differences in the set up of the machines, anything that could vary in the whole process not just in the machine. Any interaction between materials, methods, people, environment, any of those things will have to be thought of as part of the process.

To sum up, there is an unexpected complexity in understanding process control. While some facets might be judged as "standard" statistics (the notions of mean and variation for example), other facets emerge in the interaction between artefacts and tools, routines, targets, ambient environment and working practices. Clearly statisticians have been engaged in teaching SPC techniques and have addressed some of the problems. Yet on the whole, statisticians have focussed their efforts on managers and engineers. Many of this group are paying lip service to improvement methodologies without understanding what they are doing, according to Greenfield (1993), who argued that process improvement would therefore not happen: “without understanding they will surely fail as surely as they would if they had never heard of TQM (total quality management)" (p. 291). Greenfield has asserted that statisticians must make an effort to ensure their techniques are not obscured by unnecessary jargon. We are seeking to take a further step in this direction. In one way, our task may look simple, as the ideas underlying SPC are not hugely sophisticated. But from another perspective, our task is complex, as our audience might have rather limited mathematical – let alone statistical – background.

We now view the same issue of developing SPC literacy, this time from the perspective of company managers who have introduced SPC methodologies.

The Company Manager

The major impression we obtained from interviewing one company manager, familiar with the SPC methodology, was his rejection of the SPC chart as a mathematical object. M4 was adamant that there was little, if any need for anybody to understand the charts in a statistical way – he certainly did not and did not see that others needed to, since there were simple 'rules' that could be applied to read the charts:

> M4: With, for example SPC's, we'll pick on that because that's the normal format we're looking at, that obviously as you all know there's some statistics behind them – and that totally loses me never mind others. However in another sense SPC's are very simple and you only need to have a fair grasp … just a very small number of key principles. You don't need to understand all about standard deviations and how they control them, it's where they end up being where they are.

> So basically what we're providing people with is this format and the basic understanding that they need to know that there's this difference between common cause variation and special variation. If your data points appear outside of the limits you've got a big signal that you need to investigate. The concept that you need to be looking at … the width of your control limits, the gaps between control limits and one of the aims should be to reduce the variation, reduce the gap between the limits. So they're very easy principles to grasp, you don't need a great deal of mathematical training to understand this….

These ideas (repeated several times during the interview) might be an example of the invisibility of the mathematical underpinnings of the charts. But, is it possible that SPC methods could be successful if they were simply used, as M4 wished, 'like a recipe book', with little if any attention to 'why'? If so what would be the meaning assigned, for example, to the lines on the charts for the average and the upper and lower controls?

> I: Does nobody ever ask you by the way what those lines are or why they are way down there (referring to the mean and upper and lower control lines)?
>
> M4: There's always one or two. Mainly it is because [name] says it is! … Generally we are expecting people to use these charts to show them … outcomes. We use them very publicly every week in our weekly review meeting to show the teams how they are doing.

Despite being convinced that one did not need to understand the basis of the SPC charts in order for them to be used effectively, a position at odds with the statisticians' perspective, M4 was also convinced that the SPC methodology was not working well in the factory. SPC charts, he maintained, were simply being used to 'track the system', proved useful in displaying data, 'to show outcomes', but were not being exploited to improve the process:

> M4: The thing is, somehow we're not utilizing SPC charts to their full potential … we're potentially not optimizing that process, the improvement opportunities on that process.
>
> … There's using and using them isn't there? They've been around for a long time, again to be perfectly honest I would say in a lot of ways we are only just starting to get to grips with really using them. In other words for a long while we've used them as a way of presenting information so instead of looking at them in accounting terms or as you say looking at a table of numbers which is dead boring and hard to read; it's much easier to look at profit development or cost development of whatever else on a SPC. But there's a whole different ball game actually using that information as a SPC is intended to be used.
>
> I: For productivity?
>
> M4: Exactly, rather than as a better way of presenting data than a table of numbers. And that's the challenge.

Yet despite this anxiety about the effectiveness of the way the company used SPC, M4 still adhered to his belief that failure to progress along the improvement agenda was not due to any lack of understanding of the 'why' underlying the charts. He and several members of the management team and been through a six-day training session and he complained that the effects had not been noticeable in terms of an upturn in productivity:

> M4: I personally don't think the issue is in so much understanding the maths and statistics of whatever behind it. I think it's more to do with commitment and follow through on making continuing improvement happen. During that training we spent a lot of time with the graph paper out drawing SPC's ourselves, so rather than just going into Winchart and typing the data in and saying "hey presto!" and it all works for you we went right behind and did it all manually. … You would like to think having done six days' training and gone into some detail on this that a few weeks time, a few months time, the quality is going to improve, that people will just get on and do it. It doesn't happen.

M4 also did not appreciate any need for others 'lower down' in the factory to understand the charts and thought the weakness was solely in management:

> M4: It's got to be supported by management… by the whole supporting infrastructure that enables it, whatever 'it' is that you actually want to happen. … So how do we make this next leap from not just looking at the data but actually using it for continuing improvement? I honestly don't believe that the big stumbling block is because people don't understand the number side of it. I think it's that in the past we've been failing to put enough pieces of this supporting jigsaw in place to make sure that continuing improvement actually happens.

Clearly the meanings of the charts in terms of their role in process improvement were very different as expressed by the statistician and the manager[4]. M4 was, however, beginning to explore other ways to present data, as a way to catalyse the next step in improvement:

> M4: We have, I believe, been putting an over-reliance on SPC's. A while ago we said SPC's were the way to go and there's much better looking data with SPC's. What then happens is everything you look at now is on SPC's and it isn't always the best way to look at it. Especially if you are just wanting to use it to present data because people then start thinking it's just a tool to present data; it's not there to improve processes. Well they may just say it shows us how we are doing. Which is not bad but it's not really what you want because you want it to be a catalyst for action... Well some parts of the industry use EXCEL spreadsheets actually. Maybe you need multiple ways of doing it?

M4 is correct in that many factories use spreadsheets to display data, as described earlier in this chapter. His suggestion of using multiple forms to display performance on KPIs is interesting, not least because to read data presented in different ways and to appreciate the commonality underlying the representations in terms of properties of a data set (the representations highlight or suppress different aspects) demands quite a sophisticated understanding of the data set and the way the tools mediate the information about it. In fact, Wild and Pfannkuch (1999) have argued that forming and changing data representations of aspects of a system is a fundamental part of a statistical approach to understanding a system more effectively.

[4] We intend to explore how these different views might be discussed in boundary crossing scenarios we plan to organise in our ongoing project.

M4's position thus seems a little contradictory, if we note that the kind of statistical insight he is advocating – multiple ways of visualising the data – is 'mathematical', yet he has rejected the need for mathematical understanding. He sees there is a problem, otherwise he would not have spent time showing the operatives how to 'do it manually'; but he is ambivalent as to whether this is effective, or whether the utilisation of SPC charts "to their full potential" involves a mathematical component or not.

We will return to this point in the conclusions. Now, we turn to some extracts from another manager, who was in charge of process improvement and thus perhaps more closely aligned with a statistical point of view. M5 was equally concerned that the SPC charts were not being used effectively, but in contrast to M4, was concerned about problems in understanding them:

> M5: ..They (team leaders and operators) are used to seeing SPC control charts and understanding variation and not getting too carried away with it. They could probably give you a fairly brief overview of what it actually means. But whether that is 100% accurate or not I would doubt. And that's not for lack of trying because we have held countless workshops over the years. Certainly down to Team Leader level on things like SPC and things like……

So why did M5 see a role for understanding and not merely skill? He argued that it is only if readers of the charts 'engaged' with the data that they would be able to use them purposefully, to look again at them and question familiar procedures, as a first step in systemizing them:

> M5. … to make decisions, need to move from opinions based on what they *observe* hopefully, rather than basically their perception of what they think is

> happening. But it's just taking it a level beyond it, and taking it away from observations and putting it into fact, which data is – fact – unless you've measured it wrong.
>
> I: And do you find there are some people more reluctant than others to do that?
>
> M5: Definitely.
>
> I: Is that because they don't understand it? (group looks at chart – several unclear comments)
>
> M5: Yes to some extent: they need to know that that's the mean of that process [pointing to the average line in the SPC chart]. But what it tells us, it takes into account that variation as well because some weeks are better than others, that's a fact of life. So we have a control - there it is 4.1%. So what that tells us is that that process under normal common-cause variation is capable of producing up to 4.1%.

M5 also suggested that "*management* needed to take any data point down a level" (our emphasis), that is, identify the key input variables that created any output KPI, tease out the most crucial influential ones and reduce variation in them by introducing standardized procedures. It was only through these steps that process improvement could be achieved: that is by explicitly recognizing the multiple and interrelated factors that acted together to account for the common cause variation. He explained by reference to an analysis of flawed products:

> M5: Yeah say for example, what you could do is you could probably break that down and say we keep 60% of our raw material on average for silicon-related

> problems, 20% because we get spots in the adhesive, 10% is for creasing in the backing paper, or whatever … so lets say mostly it's the silicon problems. So then what you can do is you can take that down a level and say OK well … what are the things that contribute towards silicon cure, for picking up, a silicon cure? And it might be the temperatures of the silicon ovens, obviously the running speed, it might be the modifier level in the silicon, etc. etc. etc. So then you can take like real process measures then, if you look at the process measures on this type of chart I would expect to see where you've got high degrees of problems in siliconizing you would expect to see a chart for some of these variables that are not in control. So these are the ones that you can then start to do things about. Let's put a standard operating procedure in place for setting up the silicon drying oven for this particular type of product and that should be what people are interested in on an operator level.

So what were M5's understandings of the lines drawn on the SPC charts? When M5 first looked at the SPC chart shown in Figure 3, he explained the target line:

> M5: Now this is meter square per man hour. There is a target in there but its not something that we really… the target is there basically because we know these are the market demands for this year. Well, we don't know what the market demands are for the coming year but we've got as good information as we can get on that sort of situation. So in order to meet those market demands we know that we should produce at that target.

He then explained the meaning of the upper control limit and also how the control limits were calculated, and changed on the basis of previous data:

> M5: That's telling us that we are actually capable of producing up to 7700 square meters per man hour. After the change it's telling us we haven't increased that, but what we have managed to do is reduce the amount of variations so we've actually increased the minimum because before we were actually capable…. if we produced only 4800 that would have also been pretty much within the capability of the processor, that's now been increased about 5200.
>
> I: Who made that decision, to increase the lower control limit to that (pointing at graph, when LCL was raised after 26 weeks)?
>
> M5: Well the control limits are not done manually; they are based on the data.

So to appreciate the process, M5 suggested that there was a need to recognize that points outside the control limits have to be investigated as the likelihood of them occurring when the process was working as normal was just very low. He, unlike M4, was beginning to articulate some simple explanations with examples of why the 'recipes for action' had been put in place.

The team leader

We now turn to the perspective of team leaders who, as we mentioned earlier, have in recent years taken up a more important role in monitoring production, alongside their role in managing teams of operators. When asked about SPC charts, one explained:

> TL1: So the operators will input a data point – their result. And if it's within the natural processing limits there is no issue. They continue with their operation. If

> the result is outside of the processing limits then they will come to either myself or the technical officer and we work out…we do an investigation to find out what has caused us to have a rule breach [a signal] – something that is outside of the processing limits. So although it is simple for the operators the fact they input the data, the real skill comes in investigating why we've had a data point that is then sat outside of the processing limits. The critical thing for these charts to work is that we do it live. If we have a rule breach then we deal with it there and then.

It is these investigations of signals that TL1 saw as crucial to 'revealing' inefficient practices:

> TL1: It's amazing since we've been doing the graphs how many rule breaches are operator error. There was a time we used to blame when a machine stopped working or … the operators would always blame the engineers and the engineers would always blame the operators and now since we have started to take this structured approach with the charts and everything they're starting to realize how easy it is for them to make a mistake.

Thus TL1 mainly saw his role as reacting to and investigating signals, rather than seeking to reduce variation. He read the SPC charts as about identifying special causes. But, in companies moving to SPC, the TLs need to read more information from the chart and to appreciate that remaining within the control limits is a necessary but not sufficient condition for the process to be deemed satisfactory[5].

---

[5] Sometimes the process changes and moves from something that is basically stable and capable (either by design or by chance) and starts to run along, say, a higher mean (see Figure 3). If this occurs, two procedures should be set in motion. First there should be an investigation as to what happened around eight data points earlier (the probability of eight consecutive points above (or below) the mean is extremely low: so the process would have to be investigated at

We now present extracts from an interview with another team leader (TL2), who was being work-shadowed on the shop floor of a factory that was part of an international group involved in the mass production of large lengths of adhesive paper, cut to various widths. The factory produced huge roles of paper (jumbos) on which was applied a coating adhesive, with paper and 'stickiness' monitored according to the customer specifications. Speed in responding to customer demands, provided quality criteria were met and wastage reduced to a minimum, was of paramount importance. TL2 was highly experienced and had been at the factory for nearly 20 years, working his way from operator to team leader. He was in charge of one 'web' in the factory: the stream of paper that went through the factory during which time it was coated with adhesive, rolled into jumbos, slit into various widths, and finally packed off on a palette for transport to the customer. TL2 was in charge of about 10 production processes for his particular web, which he controlled and monitored by observing and manipulating dozens of dials and data windows shown on six computer screens, as illustrated in Figure 4. The web also ran past the window to his 'hut' so he could see its progress directly as well as through the data.

Insert Fig 4 here

We wanted to know TL2's interpretation of the SPC chart in Figure 3, with the KPI taken to be an output measure for his web, with the amount of product produced reported as 'square metres per manned hour'. TL2 was shown the chart and asked to describe it:

---

the start of this 'run'). Second, a 'step change' can be initiated, in which a new mean is set, and new control limits (automatically) calculated.

> TL2: So they're looking for however many square meters are going out the door in return for the manpower that they've got…. So for the five people working we're averaging from there round about 48,000 to about 78,000 square meters per manned hour. And the average is about 78 something like that, the average. Well, not the average, the optimum…efficiently at full speed constantly, I think it's 78 which is about right, that's right.

First note that TL2 was immediately reading into the SPC chart data that were not presented; for example, the five people working on his web. He read the chart (correctly from a mathematical point of view) as an average with variation, but then insisted that the Upper Control Limit was the maximum length of paper his web could produce. He made a quick calculation as to the length of paper the web could produce, adding that he knew that the machine ran at 700 meters per minute:

> TL2: UCL that's the maximum. It can't get any more than that because speed vs time, if you like. So obviously you're running at 700 meters per minute, times that by obviously 60 per hour, times 24, times 144 which is year full capacity and it won't go any higher than that unless you increase the machine speeds or alter the shift patterns ….. That line indicates 78,000 square meters per hour. *And you can't get more than that*. *It's not physically possible*. [our emphasis]. You can achieve more if you do something to the machine to make it run faster to produce more or put more hourly capacity into the working week which has 24 hours spare. Say we were running at optimum speeds with no web breaks.

TL2 went on to work out the maximum output more carefully, but then noticed something was wrong and the calculation gave a result that he knew from experience was too large. He then found reasons for the smaller number:

> TL2: But it's not is it. It should work out. 700 times 60 is 42,000 linear meters a minute, times 2 is 84 square, times 144 = 12.096 million a week. But then again you've got your start-up, shut-downs, clean-downs and…obviously. That'll be why 84 is your optimum. Your upper control limit is 78 because they've obviously got to take out your wash-downs, clean-downs…

It was hard to follow TL2's calculations as there was so much embedded knowledge there (for example, the 42,000 is multiplied by 2 to obtain the area of the paper since the paper comes in 2 meter widths). Also the shop floor is not conducive to close questioning. Suffice to say that TL2's interpretations of the control limits were largely deterministic, based on an assumption that they were determined only by the speed of the web 'minus special causes', rather than as a property of the data set. TL2 did however have a good grasp of common cause variation:

> I: But why is there all this sort of going up and down like this?
>
> TL2: Lots of reasons. It could be lack of orders, lack of raw materials, machine breakdowns, web breaks. Like I said, a lot of things. It could be a process problem, a quality problem, raw materials, lack of work. It could be all sorts, you know.

Clearly TL2 was unsurprised by the common cause variation as displayed in the chart and put it down to a range of multiple and interconnected factors – just as the managers did – but based this recognition on his knowledge of the data and the whole work process rather than of statistics. TL2 could also explain why the data stayed high over some weeks, by reference to the situation – in this case the particular product on the web:

> TL2: We're probably running the same adhesive with massive runs of the same material with a less amount of changes and breakdowns.
>
> I: That causes why we've got more output?
>
> TL2: The less change you've got the less waste you've got the more you produce.

For a dip outside the control limit, a signal, TL2 also could immediately provide a simple causal explanation.

> TL2: A particular mechanical breakdown…tension control problems, so the web must have broken...

Like TL1, he saw the main function of the charts as monitoring production and describing the reasons behind such signals, not about explicitly distinguishing the multiple factors in common cause variation mentioned earlier, and reducing it. However, TL2 did talk about his work as part of a process improvement team on waste, where again, SPC charts were being used, again not in his terms to reduce variation, but to bring down the output of waste:

> TL2: To get a tighter grip and lower, you know, if it's waste. You want to get it down, don't you?

He explained – just like M5—the way the team set out to identify the main factors in waste, one of which had been at the point of changing materials:

TL2: It's because basically the machine when it's stopped, it's not doing anything, is it. You tick it over it's doing 10 meters a minute. All that 10 meters a minute for however long you're ticking it over is wasting the relay because, it's just waste. Because you're not putting glue on it. We're putting silicon on but we're not putting glue adhesive on and we can't sell it to anybody. The machine needs to be running at a certain speed, tensions and certain coat weights in which…to make the pressure sensitive labels work correctly.

I: Oh I see.

TL2: See for 10 meters a minute it's no use. So what we'd need to do is to ramp the machine up to say 300 meters a minute or whatever speed it runs at with everything fully in place. And when, after a minute or so, when the adhesive goes through to the reeling or the laminator then we flag it and anything below that flag is waste. So we need to try and do things better, quicker and reduce waste.

I: So you've got to try and do this takeover faster.

TL2: Exactly.

I: How will you go about trying to think how to do that?

TL2: There's certain ways, and certain reasons why realistically what you'd want to do is start the machine up, put everything in place, i.e. bring your ovens down and then go to run and get up to optimum speed the quicker you can. There's certain things you can do like make the ovens get up to temperature quicker, get

> the technicians involved and make the machine run faster because it only runs at certain speeds.

When talking about process improvement, TL2 talked about actual working procedures: reducing the speed of the machine, increasing oven temperature more quickly, rather than considering 'abstract' data. He was aware that the different factors interacted on the shop floor but was unlikely to have been able to represent how they interacted with any precision or indeed quantify and compare the interactions. TL2 was also clear that although he saw the benefits of moving to more standardized practices, his knowledge of the conditions for his web to work could *not* be codified; there were just too many interacting variables to take into account to decide the best setting for any process:

> TL2: Well, we run obviously different grades of paper, well I mean grades – thicknesses of papers, 59 gram, 78 gram, 88 gram, so obviously the reason that we're heating that up, putting temperature into the paper is to cure the silicon. The thicker the paper the more temperature you need. Also you've got to take into consideration the machine speed. The faster you run the more temperature you require so… It's not just that, that's just a very, very small fraction of the whole process. You've got to remember things like that for various different products – we have 65 different products – I'm not saying we have 65 different settings, but they do vary. And in this product range we do like laid flat, reverse curl, backprintings, synthetics, you know.

TL2 also summarised all the reasons for faults, again arguing that there were too many to record:

> TL2: It's the whole…there's quite an array of things, like you know, there's all sorts of things you could be…the machine could just be ticking over and the paper could just break because you get a raw material fault. If you look at a faulty machine overall there's thousands of variables and all of these variables can have an impact on things that matter.

All this expert knowledge of the process was in his head and could not be written down:

> TL2: It's all in my head now.
>
> I: Don't you think that's odd? I mean, you've got all these values here and you have to know that the oven temperature is such and such…
>
> TL2: Yeah, we have discussed this at team leader's meetings. And they would love us to document it, but because you know, I say there's not a forever changing process but it does change and it also changes I believe at the norm for ambient temperature. In the summer, we don't need the ovens so high in the summer because the ambient temperature has an affect upon the product, believe it or not. So you could write down a whole load of settings for a load of everything, but it could change inside a month.

So, in summary, TL2 – like TL1 – tended to read the SPC charts as a series of causes that were not quantified, that interacted over time, again in ways that were not measured or even thought to be measurable: "went up as no breakdowns, good temperature and raw materials … went down as had an operator error or bad release results or silicone ovens working erratically". The control limits provided a demarcation for signals, but the role of the lines themselves and how they were derived remained unclear. Why were particular data points signals? The team leaders could certainly interpret common variation on the basis of their knowledge of the whole work process. This knowledge was shared with management to help explain variation, but the next step, to control it, may need some way of aligning meanings with a statistical view.

Towards a theoretical framework

This chapter started from the assumption that the advent of digital technology as a pervasive force at work necessitates a reappraisal of the nature of mathematical knowledge used at work and how it is communicated. We have sought to illustrate how the models of the workplace, particularly in relation to data modelling and statistical process control, play a range of roles for the different actors in the workplace activity system. We now try to draw these ideas together with a theoretical framework with which to interpret them.

Questions that we began to isolate during our case studies of the use of SPC charts were: What are the meanings ascribed to the line indicating the historical mean of a process? How are signals interpreted? How is a signal conceptualised in relation to the whole process? How are the upper and lower control limits on the SPC chart interpreted and what are the meanings ascribed to the variation within and outside these limits, especially in relation to any targets or customer specifications?

The question of control limits has turned out to be suggestive of a more general phenomenon. For a company to improve, it must ensure that the control limits are within the customer specification but in times of tight profit margins also that variation is reduced, that is, the gap between the control limits is narrowed. For some people in the activity system, but not all, this meant that at some level, the mathematical object represented by the gap between the control limits has to be understood. We say *mathematical* in the sense that there is no concrete instantiation of this gap (or indeed, of either control limit) in the process itself. Like any statistic, say, the mean, this gap is an abstraction, with no straightforward referent in the work system. In this respect, therefore, it is relevant to ask who, among those viewing and working with the charts, needs to view the control limits as mathematical objects, and for what purposes?

Consider the manager who noted that certain phenomena instantiated in the chart would necessitate taking things 'down a level'. Regardless of whether his metaphor of up/down is the right one (or in the right direction) his meaning is clear. The chart's structure would point to complexities and/or interactions between variables, and this would necessitate closer investigation of individual values – either readings on other charts, or readouts of individual machines reported in the form of abstracted data. The operators and the team leaders need, similarly, to look at particular elements of the process – down a level from the aggregated structure of the chart. However, the information they use to effect this level change tends not to be abstracted data, but actual factors they know affect the work process.

In contrast, we conclude that managers have to conceive of variation as a factor to be manipulated, and they necessarily have to see mean and variation as inseparable, i.e. they have to see the charts as mathematical – and many may not. In order to appreciate the sources of common cause variation, some of the managers see the need to analyse the multiple factors that give rise to it and then seek ways to reduce it. Others, as we saw, regard their role more as 'following recipes'. In terms of investigating special causes, this appears a reasonable strategy – they are, after all, deterministic. Where such a management approach falls down is in seeking to reduce common cause variation, which requires engagement with the data and distinguishing its interconnected components. Team leaders, on the other hand, appreciate the 'system on the ground' although they tend to see the charts merely as a way of displaying information – which does the job fine! What is lost is the stochastic dimension, and the meaning of the mean in relation to variation that depends on the historical data set.

Managers in charge of process improvement do, of course, appreciate the complexity and stochastic nature of charts. They understand the importance of the interaction of factors leading to common cause variation *and* the importance of interactions *between* KPIs, although they may not have tools to quantify random variation within *interacting* variables. This latter appreciation requires quantification and analysis at different levels of the process and represents a multifaceted, multileveled TmL. Yet we have seen that there are managers who reject the need for more quantification and the use of statistical thinking throughout the workforce, point to the importance of management, but are dissatisfied with outcomes of process improvement.

Making sense of this ambivalence is an important aspect of the problem we seek to address. In this respect, the role of the chart as a boundary object is particularly relevant, as it draws attention not only to what is different between the various communities in making sense of them, but also to what is the same. This sameness derives precisely from the fact that the charts are abstractions. For one group, the team leaders and operators, the charts are essentially 'triggers' for pinpointing (usually) human error; for another, (some of) the managers and statisticians, they are a source of data for further analysis: they, unlike the operators, are aware that remaining within the control limits is a necessary but not sufficient condition for the process to be deemed satisfactory. The fact that the charts are mathematical abstractions from the production process allows them to become shared territory. At the same time, the kinds of knowledge about the objects and relationships represented in the charts, what the variants and invariants are, the abstractions derived; differ among the actors who use them.

The roles of the charts as boundary objects gives rise to particular forms of knowledge that characterise their use, deriving from their status as abstractions from the work process. Unsurprisingly, we encountered nobody (except perhaps the statistician) who required an explicitly mathematical view of the charts, in terms of statistical theory. We did, however, find persons at all levels, who had derived ways of talking about and conceptualizing what they thought the chart represented, and a language with which to express them. We know, for example, that team leaders now have to come and present – in the words of one manager – "a lot of analytical data about what happened in their shift": they have to know more of the system (devolution) and articulate what they know (communication). These data are not a narrative of what happened, but a situated abstraction of what happened, in the sense that the quantitative readouts of the statistics and charts – encapsulating all manner of relationships between the variables concerned – are one level removed from the actual operation of the machines. We do not yet have enough data to characterise this form of quasi-mathematical expression explicitly, but we know from one of the managers that it is necessary (recall the team leader who was "not able to interpret [data], to transform it into a trend or a problem solving analysis").

The different ways to think about the charts as boundary objects accounts for the ambivalence of the managers themselves as to the role of the chart. M4's apparently contradictory view, particularly, in which he simultaneously asserts there is no need for a mathematical view of the chart, and yet works to induct his trainees to produce statistics by hand or wants them to cope with multiple representations of data sets, reflects this ambivalence. However, the question still remains as to how far is it important for the workforce *as a whole* to engage in statistical thinking to make effective decisions on process improvement.

Our initial thoughts are that negotiation of meanings of the SPC chart may provide a fruitful way forward in providing a window on this question and the work process as a whole. We describe a telling example. We know that SPC software in common use eliminates special causes from any calculation of the mean and control limits as they are not part of the stable process. One manager we interviewed was clearly confused about this important aspect of the charts. After demonstrating the software that generated the charts, pointing out the upper and lower control limits and the special cause points which operators were expected to annotate with reasons, she added "it (the special cause point) is taken out of the calculations of mean and upper and lower controls. The process automatically re-calculates without the points outside limits … not sure why … I suppose as we want to make sure we reduce the variation in process." She is right of course: it would certainly reduce the variation but not the common cause variation which is the aim of the whole SPC initiative!

Here is the nub of the issue. This manager, and perhaps the workers more generally, need to come to see data sets both as a concrete representation of what the company produces, and as an abstract representation that identifies trends with both historical and predictive significance, which points to strategies to reduce variation as well as special causes, and simultaneously provides theoretical (mathematical) and pragmatic indicators. The distinction between parts of the workforce is *not* as clear-cut as we may have thought: it simply does not seem to be the case that one sub-community "needs" one view, while another "needs" a different one. Instead, it is a question of balance, in which the mathematical element of the charts comes into and out of focus while in use for different purposes, different sectors of the communities, and at different times. We are led to concur with Greenfield's assertion that "understanding" of management is necessary for the success of SPC techniques, and we may go further in suggesting that those who are being managed need some elements of understanding as well.

We conclude by revisiting the overview of modern production processes with which we began. The unidirectionality that Reich describes – from symbol analysts downward – is no longer the only dynamic: increasingly, communication is a two-way process between those who develop the symbolic frameworks that drive machines and systems, and those that operate within them (for a recent discussion of new configurations of work, see Engeström, 2004). We should stress that *nobody* in the factories we visited could be described as a symbol analyst: they were firmly behind the scenes and we have no data even to assess how many there are. Neither, incidentally, do we know how many levels there are between our senior managers and the symbol analysts. It is, entirely possible (probable even) that those responsible for installing the computer systems in the factories were using commercially-available software, or at most, configured versions of off-the-shelf systems. Somewhere along the line there are symbol analysts, but it is quite likely that not one of them has any first-hand (or any) knowledge of the factories into which the system has been adopted.

This point is not only of passing interest. It means that in terms of the techno-mathematical knowledge of the actors in the workplace, we may be dealing with a few slices of a highly complex hierarchy, rather than two or three main delineations. Reich's more coarsegrained classification of the workplace is entirely appropriate for his economic and social perspective, but it may not be adequate for our purposes. The implication is that the knowledge that characterises the boundaries across which our team-leaders, managers and statisticians are communicating, may be different more in nuance than in real substance.

Nevertheless, some real differences between elements of the workforce are emerging. There are – as we have seen – surprising differences in the ways that the different sub-communities who engage with, for example, the SPC charts, use them, and the apparent meanings of statistical variation that are shaped by them. These meanings tend to mask the complexity of the system, and render invisible the interactions between a wide range of variables that become 'wrapped up' in the data points – only to be opened up by considering all the day-to-day specificity of the shop floor. Thus the work of the floor must be made visible as a first step in appreciating variation. The formal and quantitative indicators of work no longer 'silence' workers, as Star and Strauss put it, but require workers for their interpretation. Our conjecture is that it is explicitly recognizing these very facets and observing how they interact that could be the key to aligning meanings across communities.

Acknowledgements

We would like to thank our colleagues, Dr Phillip Kent and Dr Arthur Bakker for their contributions and comments on an earlier draft of this paper.