MATHEMATICAL EPISTEMOLOGIES AT WORK

Richard Noss

School of Mathematics, Science and Technology

Institute of Education, University of London

In this paper, I draw together a corpus of findings derived from two sources: studies of students using computers to learn mathematics, and research into the use of mathematics in professional practice. Using this as a basis, I map some elements of a theoretical framework for understanding the nature of mathematical knowledge in use, and how it is conceptualised by practitioners. I then draw some provisional implications for a set of design principles for activity systems aimed at fostering mathematical learning. I propose a central role for digital technologies, both in assisting understanding of the genetic construction of knowledge, and in developing learning environments that instantiate the design principles.

In mathematical terms, there is a celebrated tension between forms of discourse and cognition that are delicately tuned to cultural practices, and those that are focused explicitly on mathematics *per se,* recognisable by its symbolic forms and epistemological structures. This tension parallels (and is perhaps derived from) the epistemological duality of mathematical thought as both tool and object, simultaneously as a component of pragmatic activity and theoretical endeavour.

The preparation of this paper has afforded an opportunity to reflect retrospectively on this duality, and on a corpus of research in which I and my colleagues have been involved, spanning a variety of sub-fields and a couple of decades. I hope it is not too fanciful to impose upon this work a narrative that was not necessarily evident to any of us while we were engaged upon it. Here is a first outline of that narrative.

I will begin with a pervasive finding that arises from investigations with (mainly young) people expressing mathematical ideas with computers. These studies led to a series of thoughts concerning the generation of mathematical meanings that nagged away until the early nineteen nineties, when Celia Hoyles and myself began to formulate a theoretical framework for describing the phenomena we encountered. Shortly after this, we had the opportunity to work in a variety of settings with the broad common aim of elaborating the mathematics used in working practices. I will illustrate how these studies began to throw light on some fundamental questions, particularly concerning the nature of mathematical practices, and encouraged us to investigate further the problem of mathematical meaning from both cognitive and sociocultural perspectives. This effort has led to some general principles about the design of mathematical activity systems for learning, and in particular, the rather special role that digital technologies may play within them. Thus, perhaps fittingly but probably over-ambitiously, I will conclude where I began, with the assertion that digital technologies can play an unusually powerful role in helping us to understand and reshape the nature of mathematical sense-making.

Before I start, I would like to make two general observations. The first concerns my wish to consider both cognitive and social dimensions. To steer a course between these two approaches is not easy, not least because proponents of each often ignore the work of the other, or denounce as mere eclecticism any attempt at synthesis (there are important exceptions to this: see for example, Cobb and Bowers 1999; Kieran, Foreman & Sfard, 2002). One organising idea for thinking about this apparent dichotomy has been suggested by Andy diSessa (personal communication) who distinguishes between phenomena that are *distally* and *proximally* social. Much of what I have to say comes from a recognition that many phenomena concerned with mathematical meaning are proximally social, in that they manifestly involve social and cultural relations between people and within communities. But I also recognise that many facets of human thought are only distally social; while it is true that what I think, and the techniques I use for thinking and communicating are shaped both socially and culturally, I think in ways that are structured by my personal cognitive history at least as strongly as by the socio-cultural relationships in which I find myself embedded. No attempt to understand how mathematics is learned by human beings can afford to ignore this, essentially cognitive element, any more than it can afford to ignore the social and cultural relations in which cognitive activity is embedded. Thus, in what follows I hope to illustrate, not only that such a perspective need not necessarily lapse into eclecticism, but rather that the coordination of the two approaches provides a possible and even necessary methodological stance.

The second observation concerns the title of this paper. I recognise that it is bad form to tell a joke and then explain it. Forgive me then, if I explain the *double entendre* in the title. I want to talk about mathematical epistemology as it is found *in work*, to understand how mathematics is used, and how it is conceived by participants in their cultural practices. But I also want to talk about mathematical epistemology as a crucial element *at work* in learning situations; how we, as mathematics education researchers can develop, not just new approaches to teaching, but new mathematical epistemologies, that is more learnable and, at least for all but the few, more expressive.

INSIGHTS FROM OBSERVATIONS OF ACTIVITIES WITH COMPUTERS

Over some two decades, Celia Hoyles and I have engaged in studies of children and adults interacting with computational systems designed to afford mathematical expression. Throughout this time, we have noticed an interesting phenomenon, which we can simplistically characterise as follows: learners are often able to express themselves in terms that might be considered abstract, yet which seem to be bound tightly into the tools and symbols of the computational world. Learners can, in other words, say and do things with suitably-designed systems that they may be unable to say or do without them, and they can often do so in ways that are interestingly different from conventional means.

I would like to elaborate two points. The first centres on the ways that learners use technology to shape their mathematical expression, how some elements of the invariant relationships between the given objects are identified and related within the symbolic discourse of the environment. In the sense that these invariant relationships remain articulated only within activity and interaction using the notational system of the virtual world, they might not be said to constitute a formal abstraction. But to the extent that they become transformed into something coherent, reusable and general, it does make sense to consider such activity as involving an abstraction of some kind. For further elaboration of this argument, in the context of stochastic thinking, see Pratt, 1998; Pratt & Noss, in press: for a study in relation to students' conceptualisations of non-euclidean geometry from a similar perspective, see Stevenson, 1996; Stevenson & Noss, 1999: and for a recent study on 12 year-olds' understandings of symmetry and reflection, see Healy, 2002.

The second aspect is related to the first, and concerns differential performance. Put bluntly, children who may be apparently unable to express *any* relationships about their figures with pencil and paper, are able to express them quite adequately (and sometimes quite elegantly) with the computer.

Reports of differential performance depending on context are commonplace. There are consistent and widely-reported findings concerning the differential performance between adults carrying out tasks in everyday settings, and when given written assessments. For example Scribner's (1985) study of the dairy industry, Lave, Murtaugh and de la Rocha's (1988) investigation of weight-watchers, and the seminal work of Saxe (1991), and Nunes, Schliemann and Carraher (1993) on street vendors, have all shown convincingly that people who are error-prone in tests are mostly error-free in familiar practical contexts, and that there is a major disjuncture between the strategies used in the two settings. More generally, and especially since the work of Jean Lave and Etienne Wenger (1991) and others in broader anthropological contexts, we may more or less take for granted the situated view of knowledge genesis. A key insight is that people construct solutions in the course of action, and that these solutions are structured by activity. In the supermarket, for example, Lave illustrates how people avoid doing what might be classified as school mathematics not because it is too hard, but because the practice of supermarket shopping carries with it its own discourse, and its own mechanisms for meaning-making. One point that is often missed, however, is that we cannot conclude that there is nothing that passes for mathematical in shoppers' activities. The point is that when shoppers *do* use mathematics, it is supermarket mathematics, a mathematics made possible through the resources of the setting.

Since these studies, the situated cognition perspective has become ubiquitous. In its extreme version, it states that 'every cognitive act must be viewed as a specific response to a specific set of circumstances' (Resnick, 1991, p. 2). However, such arguments – compelling as they are – face researchers of mathematical learning with a number of seemingly intractable difficulties. If mathematics cannot be regarded as a decontextualised resource to be learned and then mapped onto settings, if it can only be defined in relation to specific situations, then we seem to have come close to distilling the mathematical essence out of mathematical thought.

As a way out of this *cul-de-sac,* Hoyles and myself proposed, some ten or so years ago, the idea of *situated abstraction* (a first attempt is in Hoyles & Noss, 1992) as a tool to aid in understanding how learners construct mathematical ideas by drawing on the material and discursive components of a particular setting (other attempts have been made by Cheng and Holyoak, 1985; Cheng, Holyoak, Nesbitt and Oliver, 1986; Nunes *et al.,* 1993). Situated abstraction seeks to describe metaphorically how a conceptualisation of mathematical knowledge can be both situated and abstract. It may be finely tuned to its constructive genesis – how it is learned, how it is discussed and communicated – and to its use in a cultural practice, yet can simultaneously retain mathematical invariants abstracted within that community of practice.

The idea of abstraction as a conceptualisation or a piece of knowledge lying in a separate realm from action, tools, language or indeed from any external referential sign system, is important from a perspective of mathematical discourse, since mathematical discourse is normally conceived as self-contained: it forms part of a system that has its own objects and its own rules for transforming them (see Piaget, 2000). This characteristic of formal mathematical abstraction is central to its utility: situated abstraction does not seek to challenge that utility, but questions whether mathematical abstractions can ever be fully separated from the context of their construction or application (see also Wilensky, 1991). Our broader hope is that the idea of situated abstraction will contribute to a theory of how mathematical knowledge is used or 'transferred' across settings (for other contributions to this emerging theory, see for example, Carraher & Schliemann, 2002; Sfard, 2002; Nemirovsky, in press; Hershkowitz, Schwarz and Dreyfus, 2001).

At the point in which Hoyles and myself began to formulate these ideas, they were essentially hypotheses, based only on data derived from children and adults engaged in computationally expressive media. Fortunately, we were subsequently able to test these ideas in studies of mathematics in work, affording the opportunity to focus on the situativity of mathematical meanings by investigating their use rather than their genesis.

INTO THE WORKPLACE

While it is clear that persons studied in their communities yield rich and useful data which describes what they do, it remains desirable to locate and elucidate the mathematical knowledge that they know. To achieve this aim, our group in London has employed ethnographic and interview data to capture meanings created *in situ* and the dialectical relationship of these meanings with mathematical expression on the one hand and professional expertise on the other. This has involved Celia Hoyles, Stefano Pozzi and myself in a series of studies with investment bank employees, paediatric nurses and commercial pilots; more recently, Phillip Kent and myself have been working with a group of structural engineers. The professional groups differ in substantial ways, but there are similarities: in the explicitness of their mathematical training, and in their intolerance – to a greater or lesser extent – of errors. We have developed a map of mathematical workplace activities comprised of documentary analysis, interviews with senior staff in each profession, general and task-based interviews with practitioner volunteers, and ethnographic observation of these subjects in the workplace.

I will now try to summarise some of the outcomes of this research. I will do so by sketching five vignettes, chosen to illustrate the outline of the theoretical position I wish to advance in the form of a set of provisional 'results'. The text of each vignette is based on the relevant co-authored papers that are referenced within it.

Vignette 1: The epistemological fragmentation of the workplace. The first vignette is drawn from a study conducted with a group of bank employees, part of which attempted better to understand the bankers' ways of thinking about quantitative data (see Noss and Hoyles, 1996a).

In responding to tasks involving the interpretation of functional relationships represented as graphs, the responses of the bank employees were surprisingly uniform. Most identified graphs as a visual display of numbers, as a pictorial representation of underlying data rather than as a functional relationship, and as an indication of a trend in data that allowed prediction. Where we saw graphs as a medium for expressing relationships (e.g. between quantity and time) bankers saw a display of data.

The origins of this epistemological diversity are almost certainly to be found in the tools of the system in which the bankers operate. On each employee's desk were several computers. Some, the traders and the operations staff, had three or four. On all but a very few screens, there were columns of data, graphs, and more columns of data: in every sense, graphs *were* pictures of numbers, rather than graphical representations of a functional relationship.

This epistemological standpoint with respect to graphical representations can, it seems, be thought of as the graphical face of a fragmented knowledge structure that characterises the practice of investment banking. We encountered departments specialising in the finest detail on one financial instrument, sharing a common wall but no common language with another – essentially similar – department. Of course, similarity is in the eyes of the beholder: while we might view, say, Nominal

Certificates of Deposit and Treasury Bills as flavours of similar financial instruments sharing the same (or nearly the same) mathematical structure, the bankers saw finely-tuned pragmatic knowledge and strategies, and a discourse that served to reinforce the differences between them.

> *Result 1: There is an epistemological fragmentation of the knowledge structure of the workplace that shapes, and is shaped by, the discourse of the working practice. Strategies are finely tuned to the pragmatic demands of work activities, with little tendency to strive for a theoretical orientation involving generality or appreciation of unifying models.*

Vignette 2: The role of artefacts and tools. The idea that people think and act within sociocultural contexts which are mediated by cultural tools is now commonplace. The work of Vygotsky, Luria and Leont'ev, indeed the entire corpus of work on activity theory, offers compelling evidence that individual and social acts of problem solving are contingent upon structuring resources, involving a range of artefacts such as notational systems, physical and computational tools, and work protocols (Gagliardi 1990).

Workplace settings are, naturally enough, littered with artefacts. These artefacts are, for the most part, a simple expression of work protocols, so that in routine use – and the overwhelming majority of time in working practices is spent on routine – the structure of the artefact is hidden from view. For example, in one study on a hospital ward (Pozzi, Noss & Hoyles, 1998), we found that a seemingly straightforward artefact like a fluid balance chart, contained within it the crystallised activity (Leont'ev, 1978; see also, Wertsch 1985) of the hospital community, shaping in complex – but unnoticed – ways the actions and discourse of those using it. A central part of this crystallised activity was a mathematical model of essential variables and relationships embedded in the activity: evidence for both the complexity and the invisibility of this mathematical model was gained by observing the ambiguity and uncertainty felt by a newcomer to the paediatric ward, as well as the extreme difficulty faced by the old-timers in communicating to her the structure that they had come to take for granted.

The arrival of the newcomer on the ward served to trigger a 'breakdown' or decision point within routine practice, a situation in which the models underpinning artefacts and the representational infrastructures on which their use depends, rise to the surface, and become open for inspection and negotiation by participants (and observation by researchers). That this model is normally hidden should cause no surprise: we have already noted that the purpose of an artefact is to facilitate the pragmatic activities of the workplace, not to learn mathematics or to gain insight into underlying models. Nevertheless, when breakdowns do occur, invisible relationships buried in artefacts do not suffice and there is a need for the community to understand at least some of the workings of the models, to examine their strengths and limitations, and to scrutinise the results of the mathematical labour congealed within them (see Hall, 1998, for a similar finding). At least in breakdown situations, we are abruptly made aware of circumstances that require more than mere procedural routine and the learning of work protocols, but systemic interpretation – the individual is required to make sense of what she does within the broader socio-technical system.

Result 2: Tools and artefacts shape activities and thought in ways that only become visible at times of breakdowns to routine. In disruptions to routine, individuals need to develop a broader interpretative view of the model that underpins their routine practice.

Vignette 3: The anchoring of mathematical meanings in practice. It will help to focus on a specific knowledge domain: I will turn to one of the most widely-researched topics in the field, ratio and proportion. Researchers on proportional reasoning in school and the workplace have distinguished two ubiquitous classes of strategies for making proportional calculations, *functional* (across measure) and *scalar* (within measure): see Vergnaud (1983) for a thorough analysis. Nunes, Schliemann and Carraher (1993), have suggested that scalar strategies offer a mechanism for holding on to situational meaning by keeping only one measure in view. By way of contrast, functional strategies tend to be seen as semantically sparse manipulations of numerical quantities *per se*. It appears that this difference is the reason why people tend to prefer scalar strategies, even when it results in a more computationally awkward calculation, and is the crux of the counterposition in the literature of scalar and functional approaches, in that the privileging of the former has arisen from the apparent necessity in the latter to relinquish meaning in the form of a situational referent. Nunes, Schleimman and Carraher concluded that scalar approaches are drawn from experiences in everyday situations, are more flexible and generalisable than easily forgotten algorithmic approaches, and, most relevant here, allow people to preserve the meaning of the situation by keeping variables separate and not calculating across measures.

Let us see how robust this finding is in the case of a group of paediatric nurses who are similarly expert in their field, but who have had years of school mathematical education as well as professional training. During ethnographic observations and interviews, we noticed that while all the nurses' drug calculations were carried out correctly (unsurprisingly, in written tests, nurses' responses were highly error-prone), the strategies adopted were varied and exhibited a far richer complexity than would be suggested either from our interviews or from the existing nursing literature. Of 30 episodes related to drug administration (out of a total of 250) we collected during 80 hours of observation, 26 combinations of ratios were observed with a variety of drugs, packaging and prescriptions. Of these, only four involved the nursing rule (in which all were endlessly drilled during training), while equal numbers chose scalar and functional strategies. Moreover, the nurses often opted for strategies that would, in the literature, be described as lacking in meaning.

Our interpretation of these findings is that the nurses' knowledge of concentration, that is their appreciation of the invariance of the relationship between mass and volume as evidenced in their drug calculations, was anchored in an intimate knowledge of the drug itself, as well as in the properties of familiar packaging constraints of prescribed doses. The knowledge was mutually constituted and expressed as both mathematical relation and culturally-shared situational noise – the same kind of knowledge that we encountered earlier, in the context of computer worlds and which we called situated abstraction.

> *Result 3: Knowledge is mutually constituted by a coordination produced in activity of mathematical knowledge and situational noise to form situated abstractions.*

Vignette 4: The qualitative restructuring of mathematical knowledge in activity. In a recent study, Phillip Kent and myself have been investigating the ways in which mathematical knowledge is conceived and deployed with employees of a large London-based engineering firm (see Kent & Noss, 2001; Kent & Noss, 2002). We have encountered, even with this mathematically educated group, a ubiquitous view that the majority of structural engineers do not "use mathematics" of any sophistication in their professional careers. So, while all believed that it was important for graduate engineers to have an appreciation for advanced mathematics, it is something they would rarely be expected to use:

> Once you've left university you don't use the maths you learnt there, 'squared' or 'cubed' is the most complex thing you do. For the vast majority of the engineers in this firm, an awful lot of the mathematics they were taught, I won't say learnt, doesn't surface again.

I think that this particular engineer's description of mathematics as not "surfacing" is a fortuitous one. We have seen in the case of the nurses, that mathematical knowledge becomes fused with professional knowledge, as situated abstraction, not as abstraction in its pure form. But it is, particularly for mathematically sophisticated groups such as engineers, this pure form that is readily recognisable as mathematics. Our engineer is right that mathematics does not surface; or rather, that it seldom surfaces in the form it was learned and taught. It has been transformed into something else, something at once more usable, more embedded, more noisy. Only the vestigial traces of the college mathematics taught to engineers remains in the mathematics that they actually use in activity.

The transformation in the character of mathematics appears to be not simply a quantitative one, nor merely a replacement of mathematical activity by professional expertise and experience. It represents a qualitative, epistemological and cognitive restructuring of the mathematics as it becomes 'embedded' in engineering expertise[1]. I claim that engineers' conceptualisations of this restructured mathematical knowledge are legitimately considered as situated abstractions.

I will illustrate with an example. The type of qualitative thinking that characterises the use of 'feel' in the engineering design process is exemplified by the concept of *load path*, the notion that the loads acting on a structure have to "flow down into the ground" like a kind of fluid. It is a powerful, very physical concept, and extremely useful because it provides a way of thinking about a structure before any analysis is done, allowing judgements to be made about the validity of quantitative analysis of the structure.

Formal mathematical analysis, on the other hand, is based on the assumption of static equilibrium, which assumes that nothing is moving in a stable structure, an assumption that appears to conflict with the load path concept[2]. Nevertheless, load path has become a situated abstraction of stability criteria: it allows predictions of behaviour that emerge from fusing together the actual properties of the material (e.g. steel beams) with the associated (mathematically-abstracted) forces.

The relevant point is this: engineering discourse employs, in at least one important way, a kind of knowledge which is at once about mathematical relations and about substance. The idea of flow makes no sense without something to flow through – the beams and struts of everyday engineering practice. Mathematical knowledge has been transformed to the extent that even those engaged in it do not necessarily recognise its existence. This poses sharply two questions: how does the formally-learned knowledge (e.g. the engineers' knowledge of Newton's laws, or the nurses' knowledge of the nurses' rule) become transformed both cognitively and culturally, into something new and more functional within professional practice and what connection, if any, is maintained between them?

I have no data on these questions. For the moment, the key issue concerns the transformation of knowledge, the creation of new epistemologies as a characteristic part of professional expertise. Here, at least, is the explanation of the apparent invisibility of mathematical activity. Here too is a broader, more culturally oriented perspective on the hitherto individualistic notion of situated abstraction that recognises the individual's embedding in an ambient social and cultural space.

> *Result 4: As mathematical knowledge is embedded in new settings and activities, it undergoes an epistemological and cognitive transformation. What is consciously thought of as mathematics by practitioners appears to be only the visible component of a larger, transformed body of mathematics in use that takes the form of situated abstractions.*

Vignette 5: The situativity of abstraction. The final vignette will deal with the most problematic (and so far, under-researched) issue. The challenge is to test the situativity of knowledge, to assess the extent to which knowledge in the form of situated abstraction 'transfers' to new situations (or better still, to find a convincing alternative metaphor for the notion of transfer itself).

In the study of nurses, we undertook a series of task-based interviews, in which the nurses whom we had followed on the ward were faced with situations that were progressively removed from the practices we had observed, yet which retained elements of familiar situations for them (see Noss, Hoyles & Pozzi, in press). We found that when the nurses were faced by a close simulation of their practice, they displayed similar strategies to those identified in the ethnographic studies, together with a strong sense of the invariant relationship of mass and volume – an abstraction. In these cases, the nurses' reasoning was supported by a synergy of their existing (school) mathematical knowledge and their practical experience. By contrast, an analysis of the nurses' responses to a less familiar scenario, derived from a breakdown observed on the ward, showed that when the texture of nursing practice became unavailable for any reason, the mutually constitutive elements of professional and mathematical knowledge became disconnected and the situatedness of their conceptualisation was apparent.

> *Result 5 (conjecture): The noise of a situation forms a core part of a situated abstraction. When it can be called upon in a new situation (and only then?) the mathematical knowledge can be 'transferred'.*

DESIGNING FOR CHANGE

I promised at the outset to draw the implications of the work studies, and to draw some general principles about the design of mathematical practices for learning. The hypothesis is that the ways in which people reconstruct knowledge for use in work is spontaneous, in the sense of deriving from participation in the practices of the community, and, for the most part, not formally taught within the practice. That being so, we might further hypothesise that – given the effectiveness of this kind of knowledge – we might attempt to design and construct activity systems for learning that harness the features of the workplace, at least those that we perceive as constitutive of learning. For reasons of space, I can only schematically summarise the findings and outline some challenges for the design of learning (and learnable) environments that flow from them: see Table 1.

	If ...	we should design to...
1	knowledge is fragmented and strategies pragmatic	demonstrate the power of invariants
2a	knowledge is pervasively structured by artefacts and their underlying models	supply lots of Really Useful Things
2b	people need to understand the models	make things that people can see inside
3	situated abstractions are mutually constituted by mathematical knowledge & situational noise	maximise situational noise in culturally-relevant ways
4	mathematical knowledge is transformed when it structures new activities	respect the mathematical epistemologies of new representational forms
5	situated abstractions depend on noise for 'transfer'	afford construction of new situations from old ones

Table 1: some schematic challenges of the mathematics in work findings for the design of learnable environments

In our book *Windows on Mathematical Meanings: Learning, Cultures and Computers* (Noss & Hoyles, 1996b), Celia Hoyles and myself argue generally that constructing runnable models in the form of computer programs[3], affords a compelling example of a learnable mathematics, opening unique opportunities for students to interact with a formal system. In modifying or constructing a model of a system, a student must articulate rigorously its salient relationships, describing mathematical structures in a language that can be communicated, extended, and become the subject of reflection.

There are many advocates of a similar perspective (see Hoyles & Noss, in press, for a review). In a recent study, for example, Sherin (2001) proposes that programming-based representations might be easier for students to understand physics than equation-based representations, and that programming-based representations might privilege a somewhat different "intuitive vocabulary", i.e. might tap into *different* things that people 'just know'. I would add a third point: that programming affords a rich set of situated abstractions of physical relationships that I think correspond to what he calls *a physics of processes and causation* (as opposed to algebra-physics which he characterises as *a physics of balance and equilibrium*).

It is not important whether we accept Sherin's conjecture or not: in *Windows* we refer to LogoMathematics or Programming Mathematics to emphasise that it *is* a different kind of mathematics that is at issue (this is an instance of the fifth design challenge). What is important is that we recognise that the switch from one representational form to another, carries with it the possibility of a switch simultaneously in epistemology and learnability.

I would like to add one more crucial element to the consideration of design principles for learning environments: the importance of mathematical models, a proximally social issue that I briefly touched on in Vignette 2. I believe the knowledge economy has massively broadened the number of people who need to understand the system they are using: elsewhere (Noss, 1998; 2002) I elaborate a case that competence in constructing, interpreting and critiquing models has become a core part of social and professional life in the twenty-first century. Sharing, critiquing and representing models is massively under-represented in mathematics curricula, still wedded to the epistemological and pedagogical requisites of the nineteenth century rather than transforming both in the face of the demands, and computational possibilities, of the twenty-first (see Kaput, Hoyles & Noss, 2002).

I contend that manipulating, modifying, constructing and sharing computationally instantiated models of mathematical systems affords the best chance we have for designing a more learnable mathematics, and of realising the five challenges outlined in the previous section. We have recently completed a study aimed at instantiating this approach in the *Playground Project*[4], which has involved a group of researchers based in four European countries who have developed a system with which young children, aged less than 8 years old, can play, share, construct and rebuild computer games. Our goal has been to put children in the role of game designers and game programmers, rather than merely consumers of games programmed by adults.

I can, for limitations of space, only sketch an example (see Noss, 2002 for a full description). Mitchell is an eight year-old boy in an inner-city school who has been helping to design and debug the *Playground* system for about a year. He is playing a game where he controls a character called 'dusty', who shoots out flowers every time the force joystick trigger is pressed. Mitchell collects points by hitting an animated target moving vertically up the left-hand edge of the screen. He finds it very easy to play and achieve high scores since he can move his character as close to the target as he pleases.

Mitchell decides that the game would be more fun if it was competitive, so at his request, Miki, the researcher, adds another player character, this one controlled by the mouse. It is at this point that Mitchell plays a trick: demanding that she avert her eyes, he removes a piece of Miki's program, effectively disabling her mouse by restricting it only to vertical movements, while he has two-dimensional control and can get as close as he likes to the target! Mitchell made use of a surprising and difficult fact: that two-dimensional motion can be instantiated as the vector sum of horizontal and vertical components.

Think for a moment of the knowledge congealed in the innocent phrase "vector sum". Concealed in this phrase, is a taken-for-granted representational infrastructure that includes the definition of a vector, the algebraic system for combining two or more vectors, and a range of properties (e.g. scalar and vector product) that give meaning to the very idea of what a vector is and why it is a conceptually powerful generalisation of a real number. This structure is relatively complex, and is postponed with good reason until the latter stages of compulsory education, if it is taught at all. Yet the complexity is in the infrastructure, not the idea: the latter is, as Mitchell showed, rather intuitive. The point is that what is intuitive is hugely contingent on the representational infrastructure with which the intuition is expressed. In Mitchell's world, the addition of vectors is instantiated not as an algebraic relation but as a natural property of the representational system. The (object-oriented) structures of the system translated, more or less directly, into what kinds of things Mitchell could take for granted as "just so", what meanings he could derive from them, and most importantly, the ways in which he could make the ideas work *for him* in achieving his goal. In short, the change in representational infrastructure transformed not only the learnability of the mathematical knowledge, but the mathematical epistemology *at work* in the activity system.

CONCLUDING REMARKS

This last point brings me to the intention I flagged at the outset, to conclude with the notion of epistemology *at* rather than *in* work. What is the connection between the two? A key link is that the analysis of mathematics in work concerns the transformation of knowledge as it is recontextualized across settings. We have seen how a person's mathematical knowledge is not invariant across time and space; it is transformed into different guises, different epistemologies, more or less visible as mathematics. This transformation seems much more powerful than the traditional notion of "application" or "use" that is often employed as a metaphor to describe this process. If formally-taught mathematical knowledge is transformed in this way, it is at least possible that reciprocal transformations may, in the future, take place for Mitchell, and that he may come to recontextualise his piece of knowledge about vector addition in (for him) novel ways.

More generally, the sketch I have provided offers a further point of connection between cognitive and cultural perspectives. In imagining how mathematical structures can be externalised and manipulated within an appropriate symbolic or linguistic framework, it suggests how abstractions constructed within concrete situations may compensate for their lack of universality by their gain in expressiveness. When general relationships can be expressed, they can be explored and become familiar. In the process, the links with knowledge of lived-in-cultures can be maintained, rather than severed in the quest for ultimate pinnacles of abstraction.

Mitchell was immersed in a world that was, I think, every bit as concrete and real to him as the load path on the components of a bridge are to an engineer. And, like his professional counterparts, Mitchell was engaged in an activity that researchers in the field of mathematical learning may recognise as having a mathematical component, but which were to him part of the ecological system – the totality of relationships between himself and the environment, and the ways in which these were expressed and communicated.

That the Playground and mathematical epistemologies run side-by-side should not be a matter of surprise: there is, after all, no single way in which humans can conceptualise (mathematically or otherwise) their environment, even though some are socially and historically privileged within a given culture. Official, symbolic mathematics is privileged in just this way; and there are good reasons for this. But the compactness and elegance of mathematical expression does not necessarily make it equally functional for learning, and if learning is our prior goal, we would do well to think about new epistemological frameworks in which to embed the mathematics we wish our students to understand. New epistemologies mean new intuitions, new things to be built with them, and new means for combining and reconstructing them. They involve new sets of situated abstractions that are both functional and powerful. I think this is the major challenge for the design of didactical environments, to create new systems which might, I think, be justifiably described as new mathematical epistemologies at work.

ACKNOWLEDGEMENTS

The corpus of data on which this paper is built was the result of several research projects, and some ten or more years of collaboration with colleagues. I acknowledge the support of the Economic and Social Research Council UK (grants R 0222 50004 & R 0002 23420) and the EU (grant 29329). Almost all of the data presented here has been published in collaboratively-written articles, which are appropriately referenced. My thanks to Ross Adamson, Miki Grahame, Lulu Healy, Celia Hoyles, Phillip Kent, Dave Pratt and Stefano Pozzi for their collaboration in the various projects. Lulu offered helpful comments on an earlier draft; Celia helped me redraft (and stay sane) when Version 1 exceeded the page constraints by a factor of 2. *Windows on Mathematical Meanings* has provided a source for the main theoretical ideas, and in this respect – as in so many others – I would like to acknowledge the profound help and support of my co-author and collaborator, Celia Hoyles.

REFERENCES

Artigue, M. (in press). Learning Mathematics in a CAS Environment : The Genesis of a reflection about instrumentation and the dialectics between technical and conceptual work. *International Journal of Computers for Mathematical Learning.*

Carraher, D. W., & Schliemann, A. D. (2002). The transfer dilemma. *Journal of the Learning Sciences*. 11,1, pp. 1-24.

Cheng, P. W., Holyoak, K. J., Nesbitt, R. E., & Oliver, L. M. (1986). Pragmatic versus syntactic approaches to training deductive reasoning. *Cognitive Pyschology*, 18, 293-328.

Cheng, P., W., & Holyoak, K. J. (1985). Pragmatic Reasoning Schemas. *Cognitive Psychology*, 17, 391-416.

Cobb, P., and Bowers, J. (1999). Cognitive and situated learning perspectives in theory and practice, *Educational Researcher*, 28 (2), pp. 4-15.

diSessa, A. (1980). Momentum Flow as an Alternative Perspective in Elementary mechanics. Division for Study and Research in Education, MIT.

diSessa, A. (1987). Artificial Worlds and Real Experience. In: R. W. Lawler & M. Yazdani (Eds.), *Artificial Intelligence and Education. Volume One: Learning Environments and Tutoring Systems* (pp. 55-78). New Jersey: Ablex Publishing Corporation.

Gagliardi, P. (1990). *Symbols and Artifacts: Views of the Corporate Landscape*. New York: Aldine de Gruyter.

Hall, R. (1998). Following mathematical practices in design-oriented work. In C. Hoyles, C. Morgan & G. Woodhouse (Eds), *Rethinking the Mathematics Curriculum*, (pp. 29-47). London: Falmer Press.

Healy, L. (S.) (2002) Iterative Design and Comparison of Learning Systems for Reflection in Two Dimensions. Unpublished Doctoral Dissertation. Institute of Education, University of London.

Hershkowitz, R., Schwarz, B. B., & Dreyfus, T. (2001). Abstraction in context: Epistemic actions. *Journal for Research in Mathematics Education*, *32*, 195-222.

Hoyles, C. & Noss, R. (1992). Looking back and looking forward. In: Hoyles C. & Noss R. (Eds.) *Learning Mathematics and Logo.* pp. 431-468. Cambridge: MIT Press.

Hoyles, C. & Noss, R. (in press). What can digital technologies take from and bring to research in mathematics education? In: Bishop A. (Ed). *Second International Handbook of Mathematics Education*. Dordrecht: Kluwer Academic Press.

Hoyles, C., Noss R., and Pozzi, S. (2001) Proportional Reasoning in Nursing Practice. *Journal for Research in Mathematics Education* 32, 1, 4-27.

Kaput, J., Noss, R. & Hoyles, C. (2002) Developing New Notations for a Learnable Mathematics in the Computational Era. In English, L. (Ed) *Handbook of International Research in Mathematics Education.* Lawrence Erlbaum. pp. 51-75.

Kent, P. & Noss, R. (2001). The mathematical components of engineering expertise: the relationship between doing and understanding mathematics. *Proceedings of the IEE Second Annual Symposium on Engineering Education: Professional Engineering Scenarios* Volume 2, pages 39/1 - 39/7. 2001.

Kent, P. & Noss, R. (2002). The Mathematical Components of Engineering Expertise: Final Report to the ESRC. http://www.ioe.ac.uk/rnoss/MCEE

Lave, J., & Wenger, E. (1991). *Situated learning: legitimate peripheral participation.* Cambridge, Cambridge University Press.

Lave, J., M. Murtaugh, & O. de la Rocha (1985). 'The dialectic of arithmetic in grocery shopping', in B. Rogoff and J. Lave (Eds) *Everyday Cognition: Its*

Development in Social Context Cambridge, Massachusetts, Harvard University Press. 67-94.

Leont'ev, A. N. (1978). *Activity, Consciousness and Personality*. Englewood Cliffs, NJ: Prentice–Hall.

Nemirovsky, R. (in press). On Guessing the Essential Thing. In: K. Gravemeijer, R. Lehrer, B. van Oers and L. Verschaffel (Eds). *Symbolizing, Modeling and Tool Use in Mathematics.* Dordrecht, The Netherlands, Kluwer Academic Publishers.

Noss, R. (1998) New Numeracies for a technological culture. *For the Learning of Mathematics*, 18, 2, 2-12.

Noss, R. (2002) For a Learnable Mathematics in the Digital Culture. *Educational Studies in Mathematics.*

Noss R. & Hoyles, C. (1996a) The visibility of meanings: modelling the mathematics of banking. *International Journal of Computers for Mathematical Learning* 1,1, 3-31. 1996.

Noss, R. & Hoyles, C. (1996b) *Windows on Mathematical Meanings: Learning, Cultures and Computers.* Dordrecht: Kluwer.

Noss, R., Hoyles, C., and Pozzi, S. (in press) Abstraction in Expertise: A Study of Nurses' Conceptions of Concentration. *Journal for Research in Mathematics Education.*

Nunes, T., Schliemann, A. D., & Carraher, D. W. (1993). *Street mathematics and school mathematics*. Cambridge, UK: Cambridge University Press.

Piaget, J. (2000). *Studies of reflecting abstraction.* R. L. Campbell, (Ed. and Trans.). London: Taylor and Francis.

Pozzi S., Noss R,. & Hoyles C. (1998). Tools in practice, mathematics in use. *Educational Studies in Mathematics.36*, 105–122

Resnick, L. (1991). Shared Cognition: Thinking as a Social Practice. In L. Resnick, J. Levine, & S. Teasley (Eds.), *Perspectives on Socially Shared Cognition* (pp. 1-20). American Psychological Association.

Saxe, G. B. (1991), *Culture and Cognitive Development: Studies in Mathematical Understanding*, Lawrence Erlbaum Associates, Hillsdale, N.J.

Scribner, S.: 1985, 'Studying working intelligence', in B. Rogoff and J. Lave (Eds) *Everyday Cognition: Its Development in Social Context.* Cambridge, Massachusetts, Harvard University Press. 9- 40

Sfard A. (2002). There is more to discourse than meets the ears: Looking at thinking as communicating to learn more about mathematical learning. In C. Kieran, E. Foreman., and A. Sfard (Eds.), Special Issue: Bridging the individual and the social: Discursive approaches to research in mathematics education, *Educational Studies in Mathematics.* 46, 1-3, 13-57.

Sherin, B. (2001) A comparison of programming and algebraic notation as expressive languages for physics. *International Journal of Computers for Mathematical Learning.* Vol. 6, I, 1-61.

Stevenson, I.J.(1996) *The Iterative Design of a Computer-based Microworld for Non-euclidean Geometry*. Unpublished PhD thesis. University of London.

Stevenson, I. & Noss, R. (1999) Supporting the evolution of mathematical meanings: the case of non-euclidean geometry. *International Journal of Computers for Mathematical Learning*. 3, 3, 229-254.

Vergnaud, G. (1983). Multiplicative Structures. In R. Lesh & M. Landau (Eds.), *Acquisition of Mathematics Concepts and Processes* (pp. 127-174). London, New York: Academic Press.

Wilensky, U. (1991). Abstract mediations on the concrete and concrete implications for mathematics education. In I. Harel, & S. Papert (Eds.), *Constructivism* (pp. 193-204). Norwood, NJ: Ablex Publishing Corporation

NOTES

[1] Michèle Artigue (in press) has made a related point deriving from the work of Yves Chevallard.

[2] diSessa (1980) makes a compelling case that a view of force as momentum flow may more easily engage and refine (rather than deny) students' existing intuitions, and therefore present a more learnable physics than that represented by the familiar $F = ma$.

[3] Constructing a computer program no longer necessarily involves writing lines of text.

[4] The Playground project was a consortium across four countries, directed by myself and Celia Hoyles. The London team also comprised (at various times) Ross Adamson, Miki Grahame, Sarah Lowe and Dave Pratt. Ken Kahn, the author of *ToonTalk,* was a consultant to the project.