DAVE PRATT

HOW DO TEACHERS FOSTER STUDENTS' UNDERSTANDING OF PROBABILITY?

Chapter 7

1. TWO KEY IDEAS IN PROBABILITY

Probability is unusual in many respects. As a knowledge domain, it straddles mathematics in its pure abstraction, and physics, economics and indeed most sciences and social sciences because of its wide range of applicability. Equally unusually for an aspect of mathematics, it explicitly pervades our everyday lives whereas most aspects of mathematics are hidden and, although they may have fundamental impact on our lives, for the most part, we are unaware of this insidious effect (Noss, 1997). In contrast, the language of probability pervades almost everything we do: sports commentators talk about a 50/50 ball, weather forecasters announce an 80% chance of rain; health is assessed in terms of risk factors based upon probabilistic calculations. Indeed it seems probability is one of the few areas of mathematics that informs explicitly the way in which we conduct our everyday lives.

It is not uncommon for modern mathematics curricula to recognise the significance of chance and probability. As an illustrative example, I will draw upon the National Numeracy Strategy for England and Wales (DfEE, 1999) but a similar case could be based on many other national curricula. This document, often referred to as "the framework" for mathematics, sets out the teaching programmes, referenced against identified key objectives for ages 4 to 13 years. Nearly all teachers in state education in England and Wales follow this programme.

Probability becomes an explicit part of the curriculum from age 7 years on. Between ages 7 and 10, the focus is on: (i) the language of probability with some emphasis on equally likely outcomes, (ii) events

Graham A. Jones (ed.), Exploring probability in school: Challenges for teaching and learning, 182—201.

that consist of two or more outcomes, (iii) how the results from an experiment can vary and (iv) the difference between theoretical and experimental probabilities. So we see at this stage some recognition of the importance of variation, though it is limited in scope, and some elementary work on calculating probabilities.

In 2001, the corresponding teaching plans for ages 11 to 13 were published (DfEE, 2001). There is now an increased emphasis on calculating probabilities and the calculation of all possible combinations in various situations. There is also some emphasis on the estimation of probabilities from experiments. A key objective aimed only at the most able students at age 13 states, “Recognise that, with repeated trials, experimental probability tends to a limit…” (p. 283).

It is questionable that sufficient emphasis is given to randomness in terms of time in the primary phase, and to the law of large numbers in either the middle school or secondary phase in terms of the range of ability for whom this is a key objective. In my view, the curriculum sends a clear signal that the ideas behind the law of large numbers are beyond the scope of all but the highest abilities.

Furthermore, there is a notable omission from the curriculum. Mathematicians and statisticians would surely argue that the concept of distribution is central to their domain. The discussion about the comparison of theoretical and experimental probabilities, which is fostered by the curriculum, should find expression through the emergence of a notion of distribution. By limiting the experience of randomness and variation to situations for which children often already have an intuitive feel, they are not given in my view the opportunity to recognise the powerful connection between randomness, the law of large numbers and distribution.

Perhaps one reason for the limited extent to which these key concepts are addressed lies in their perceived difficulty. Teachers need to find ways of building on what children already know and to be aware of the limitations of that knowledge if they are to find pedagogic strategies that support the learning of these concepts.

2. WHAT DO CHILDREN NOT KNOW AND WHAT *DO* THEY ALREADY KNOW?

The domain of probability and chance has been the focus of a great deal of research into the errors and irrational thinking that people, not just children, seem to exhibit when making judgements of chance. The failure of our intuitions has been so well documented that it is perhaps not surprising when teachers, confronted with the difficulties faced by their children, believe probability is simply counter-intuitive. A corollary to this view could be that our mental apparatus is hard-wired in such a way that it is beyond redress through any pedagogic strategy. It is certainly worth briefly summarising that body of literature before evaluating whether the above perspective is the only defensible interpretation.

Research on What Children (and Adults) do not Know

The seminal work was carried out by Piaget and Inhelder (1951/ 1975). They noted that in order to accommodate probabilistic thinking the organism needs the capacity to recognise uncertainty and to be able to catalogue systematically all possible combinations. The latter requirement demands that probabilistic knowledge is a late development, well into the stage of formal operations. Probability theory can be seen as an invention by the organism to operationalise randomness.

Meanwhile, how are people to make judgements of chance in everyday life? Many researchers have offered descriptions of the sorts of heuristics that people use to make such judgements. The main body of literature has been provided by Kahneman and Tversky (1982) who catalogued during the 1960s and 1970s a long list of such heuristics. They provide two of the four heuristics outlined below.

The *representativeness* heuristic (Kahneman & Tversky, 1973) proposes that, in making a judgement about a chance situation, we assess how well the information in the situation matches or represents what is perceived to be the parent population. Such a heuristic breaks down in various well-documented situations. For example, a run of reds on the roulette wheel seems to generate amongst many people the judgement that the next outcome will be a black (so-called *negative* *recency*). If the next result were a black, the outcomes as a whole would certainly be more representative of the population in which there are two possible outcomes, red or black. However, the roulette wheel has no memory of what has already happened, and so, unless the roulette wheel is broken, there is no reason to suppose that a black is more likely than a red on the next roll. The representativeness heuristic is a poor means of making predictions at the roulette wheel, and indeed in many other everyday situations.

Kahneman and Tversky (Tversky & Kahneman, 1973) also describe the *availability* heuristic as another means by which we often make errors in our judgements of chance. In this heuristic, we make judgements by evoking from memory prior experiences of similar situations to try to gain a sense of how often they occur. Many parents are familiar with the dilemma of deciding whether to allow their children to make their own way to school. The dangers are self-evident. A traffic accident is just one danger that preys on the parent's mind. In making the judgement, we are likely to be heavily influenced by stories in the media or unfortunate experiences of friends or the family. The salience of these incidents is likely to persuade us to take the safe line. From a parenting perspective this may be entirely justified but as a judgement of chance it fails to account properly for the many thousands of times that children have indeed found their way to school safely. Unfortunately, according to Kahneman and Tversky, our ability to attach probabilities to events is severely distorted by the salience of evoked memories of similar past events.

In many situations we do not even recognise that we might usefully apply a stochastic model to a situation in trying to make a judgement of chance. In strategy games, for example, we are often so pleased to have carried out come successful coup that we fail to recognise that we were "lucky" in the sense that we used a strategy that most of the time would have been unsuccessful. This attention to the result rather than the probabilistic strategy has been called the *outcome* approach (Konold, 1989). The outcome approach is one reason why, when teachers ask children to make a prediction about a chance situation, the children will respond that it is impossible to say, "it's just a matter of chance". For them perhaps, all that matters is what happens in practice.

Lecoutre (1992) has reported a related phenomenon, named the *equiprobability bias*. In a fascinating study, she noted that, when children were given two tasks with identical mathematical structures, they performed far worse when the task had a random element to it. She offered her students (15 to 17 years of age) three square cards, depicting triangles on two of the cards and a square on the third. The shapes were configured in such a way that the students could generate a house shape (in two ways) or a rhombus (in one way) by placing pairs of cards together. The students were shown how the house and rhombus shapes could be constructed and were asked whether a house was more, equally or less likely than a rhombus. The same subjects were then asked a similarly structured question where two orange flavoured and one lemon flavoured candy were placed in a bag and the students were asked to compare the likelihood of choosing various combinations of pairs of candy. Finally the students were asked to compare the two experiments. The results showed the number of correct responses diminished markedly from question 1 to question 2. Lecoutre argued that the equiprobability bias was resistant to modification (even amongst individuals grounded in probability theory) but that a correct response could be induced by masking the chance element of the problem. She concluded that correct cognitive models are often available but are not spontaneously associated with the situations at hand.

In other words, children who were quite capable of identifying the possible combinations typically failed to use this information correctly when a random element was added to the task. They would tend to respond instead that it was just a matter of chance or "50/50". Lecoutre's work suggests that it is perfectly feasible to gain success by masking the random element in a task, and that our lack of comfort with randomness persists even beyond the point in our development when we are able to compute combinations.

The above account of human fallibility in making judgements of chance is depressing but it is not at all clear that the catalogue of failure necessarily implies that it is impossible to offer children productive learning experiences. Indeed the next section will strike an altogether more optimistic note.

Research on What Children (and Adults) do Know

Piaget's approach was to examine the epistemology of probabilistic knowledge from a genetic perspective and as such he was less interested in how the setting might shape such development. In contrast, teachers deal continuously with the partial knowledge of their children, and consequently teachers need guidance on how their actions, including the offering of certain types of resources, might shape children's knowledge development. Fischbein's work (1975,1996) on intuitions provides some constructs that relate to those aspects of our mental apparatus that are brought into play when we are making more immediate decisions and judgments. Fischbein found evidence that at very young ages children (even pre-schoolers) are able to intuit relative frequencies. Fischbein's thesis suggested that the weaknesses described above were the consequence of a pedagogic emphasis on the deterministic. According to Fischbein, children's early *primary* (by which he meant unschooled or untaught) intuitions fail to develop as effective *secondary* (scientifically learned or taught) intuitions because of a lack of support from the school system. However, his work does not quite reach the level of specificity that would be of fundamental help to the beleaguered teacher.

Nevertheless, Fischbein offers a more positive outlook in the sense that his work promotes the notion that new pedagogies might support the development of "better" intuitions, rather than leaving our children to develop in a state of epistemological anxiety (Wilensky, 1997).

My own research has reported children at age 10 or 11 years with well-established intuitions for randomness. Young children seem to recognize random experiments as involving the following characteristics: unpredictability, irregularity, unsteerability, and fairness (Pratt, 1998a).

1. *Unpredictability*: If the next outcome is not predictable, a child might regard the experiment as random,
2. *Irregularity*: If there is evidently no patterned sequence in prior results, a child might refer to the experiment as random,
3. *Unsteerability*: If the child is unable to exert physical control over the outcome of the phenomenon, the experiment might be seen as random, and

4. *Fairness*: If there seems to be a rough symmetry in the experiment, a child may think of the experiment as random.

I would claim that these intuitions for randomness are not so different from the expert perspective, though, whereas these four intuitions might be about as much as a 10 year old child knows, the expert's knowledge will connect these intuitions to a rich and extensive concept image (Tall & Vinner, 1981).

An expert recognises the differences between fair and random. In particular, random can be biased and so perhaps might be regarded as unfair. My perspective is that we should not dismiss the child's knowledge as a misconception to be eradicated (see Smith, diSessa, & Rochelle, 1993 for a brilliant articulation of this perspective). Instead, we should *accept* the pedagogic challenge of how to build on the child's impoverished view of randomness so that it is connected to, but not identical with, that of fairness.

From an expert perspective, the four intuitions for randomness sometimes appear self-contradictory. To a young child, a spinner, whose equal-sized sectors read 1, 2, 3, 4, 5 and 6 might be seen as fair and so random. Now consider a spinner also numbered 1 to 6 but in such a way that the 6 sector is twice the size of the others. The same child might well regard this spinner as unfair and so non-random. Both these spinners are in fact largely unpredictable and unsteerable and will generate irregular results, and so in these respects the experiment with the non-uniform spinner might have been regarded as random too. In my research, children often appeared unconcerned by such inconsistencies; they adopted whichever stance was cued by the most obvious characteristics of the situation in question. The teacher's role might be to find a way to raise these inconsistencies in the hope that the cognitive conflict somehow helps the child to begin to distinguish between fairness and randomness. Of course, the problem for the teacher might be that the child continues to ignore the conflict, which might appear to be more of a conflict for the teacher than it is for the child!

There is one final point that I wish to raise about the differences between the limited, but useful, intuitions of the 10-year-old child and the powerful expert understanding, and this is in my view the most significant issue. The four intuitions for randomness focus entirely on immediately observable aspects of the experiment. The children in my research did not in the initial interviews exhibit any awareness of the longer term aggregated properties of randomness, Yet, the crucial understanding that an expert has, and one of the key objectives identified in the first section, is that set out in the law of large numbers. The mathematically exciting property of random experiments is not so much their unpredictability in the short term but their predictability over a large number of trials, in the sense that the relative frequency of an outcome tends towards its probability.

The law of large numbers, regarded by the Numeracy Strategy for England and Wales as accessible to only the most able children, appears to be the principal aspect of randomness that distinguishes an expert understanding of randomness from that of some 10-year-old children. How might teachers support the development of intuitions for this idea (and indeed for distribution)? To begin to answer this question, I wish to summarise one piece of research involving the use of technology and to infer some pedagogic principles that might guide our teaching.

3. THE CHANCE-MAKER STUDY

In providing this synopsis of part of my own research I will illustrate what I see as a number of important pedagogical guidelines that hopefully flesh out Fischbein's precept: greater emphasis on stochastics would provide better support for children's early intuitions (1975, p.73).

My study was in fact a piece of design research (Cobb, Confrey, diSessa, Lehrer, & Schauble, 2003) in which one aim was to build a microworld, which eventually became known as *Chance-Maker*. At the same time, I aimed to gain fresh insights into how children's stochastic thinking evolved through the use of the developing computer-based tools. The findings about these children's initial intuitions for randomness have already been discussed. In this section, I wish to emphasise the emergence of new knowledge.

The Chance-Maker Microworld

In the final iteration of the research, the children were given a series of *gadgets* (Figure 1), mini-computational devices that simulate everyday random generators (a coin, a spinner, a dice and so on). An assumption was that the children would regard the normative state for such gadgets as one of being fair. In order to appreciate their understanding of chance, I needed to challenge this perspective. Hence, the gadgets were in some cases intentionally broken, in the sense that some sort of bias had been inserted into their operation. The children were asked to identify which gadgets were not working properly. The gadgets also contained a variety of tools. The children were challenged to use these tools to mend the broken gadgets. My assumption was that they would aim to make the gadgets fair but the research showed that fairness has many ways of manifesting itself.

Each of the gadgets shown in Figure 1 has a strength control. This allows the child to control how hard the coin, spinner or dice is thrown or tossed. Higher strengths make the simulation continue for a longer time period though in fact strength has no effect on outcome. Alternatively, the child can click with the mouse directly on the gadget, in which case it is triggered with the same strength as last time, allowing replications of experiments (which in fact do not necessarily generate the same outcome).

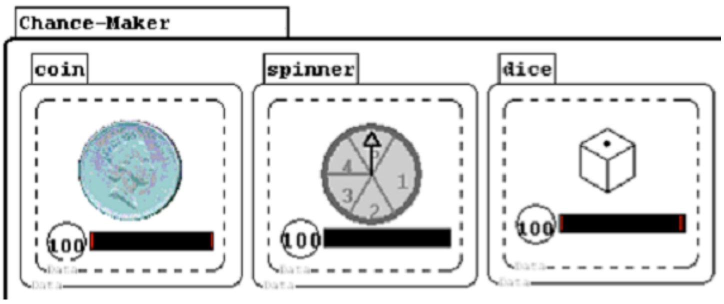

Figure 1. Three of the gadgets in the Chance-Maker microworld

When a child wishes to mend a gadget, she opens up the gadget to gain access to the mending tools. In Figure 2 the tools for the dice gadget are shown. The results are listed (in the *Results* box) and can be displayed as a pictogram (*Pict* button) or as a pie chart. (*Pie* button). Trials of an experiment can be repeated many times (in Figure 2, the *Repeat* tool is prepared for an experiment of 100 trials), usually by turning the graphics off to save time (*On/Off* button). Results will accumulate until a new experiment is begun (*New* button).

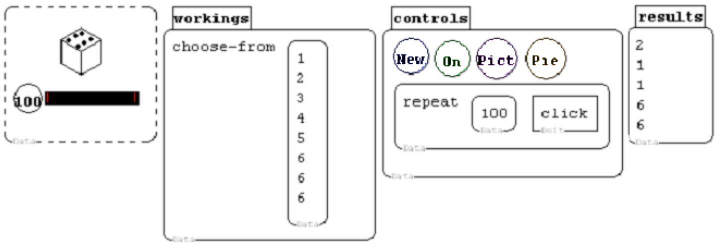

Figure 2. The main toolsin the dice gadget

The workings box shows the computational core of the gadget. In this particular case, the dice "chooses" from the list 1, 2, 3, 4, 5, 6, 6, 6. The workings box can be edited by the child to change the way it works.

Emergent Knowledge

I summarise below a typical evolution of knowledge. For my purposes here, which is to abstract the pedagogic issues, this summary is sufficient. Of course, there were variations in the ways that different children interacted with the software but this summary is typical of the children I worked with. The qualitative methodology adopted for this study does not allow claims of statistical generality, though hopefully the reader may find some resonance, which in a sense imbues the work with a degree of generality (for more detail, you may wish to refer to Pratt, 1998b, 2000).

The children interacted at top level without using the tools to ascertain which seemed in their view not to be working properly. There was much evidence to support the intuitions for randomness identified in the pre-interviews and described above.

By default, the coin gadget was unbiased. Nevertheless the children thought that they identified effects due to the strength control. The tools helped them to ascertain that there was no maintainable pattern to the results whatever conjecture they held for how the strength was impacting upon the results.

They were confused that the pie chart for the coin did not appear uniform. They experimented again with the strength control and various other features to see if they could make the pie chart display as they felt it should. They would consistently use small numbers of trials and so the pie chart was never satisfactory.

Sometimes by accident (the software accumulates results unless the *New* button is pressed), and sometimes after a researcher prompt, the children tried increasing the number of trials and found that the pie chart would then appear to be more even. This insight can be characterised as “the more the number of times we spin the spinner, the more even is its pie chart”. This articulation is, in my view, a good example of what Noss and Hoyles (1996) call a *situated abstraction*. The children have abstracted a rule for how the phenomenon behaves but the abstraction is apparently tied to its setting in so far as one can ascertain from their language.

The spinner gadget looked unfair and so the children quickly suspected it was not working properly (they felt it should be unbiased). Their attention was drawn by this unfairness to the unfairness of the workings box. However, editing the workings box so that each outcome appeared only once did not seem to solve the problem. The pie chart for example still appeared “unfair”, in the sense that the sectors were unequal.

Instead of reusing their situated abstraction from the coin gadget, they continued to use small numbers of trials and tried many configurations of the workings box, adjusting the values to compensate for discrepancies in the previous pie chart. Nevertheless, the pie was inconsistent in its appearance. This activity seemed to demonstrate the deep situatedness of the knowledge gained from working with the coin gadget.

Eventually, perhaps out of desperation, they recalled quite explicitly what they had learned about the coin gadget. They tried a greater number of trials and found at last that the pie chart now appeared to them to be fair.

Along the way they articulated the situated abstraction that “the workings box makes the pie chart fair”. However, when challenged by the researcher, some of the children realised that this would only work when the number of trials was high.

The children began working on the dice exactly as they had with the spinner. Again they ignored their previous learning. However, on this occasion they were much quicker to turn to their previous situated abstractions to explore what would happen with a uniform workings box and a high number of trials. It seemed that, by the third gadget, the new knowledge was now sufficiently reliable to be called up in preference to other ideas that they might have for how it worked.

This brief synopsis only covers some aspects of the research but nevertheless it allows me to draw out come pedagogic guidelines, which the teacher might find resonant with their own experiences in the classroom.

4. IMPLICATIONS FOR PEDAGOGY

Purpose and Utility

Teachers are confronted with what Janet Ainley and I have called the *planning paradox* (Ainley & Pratt, 2002). The Numeracy Strategy in England and Wales is, like most curricula, set out in terms of teaching objectives, based directly on mathematical concepts or skills. Curricula for middle schools (say, 11 to 15 years of age) across the world tend to be presented in a similar fashion. However, we claim that if teachers plan from teaching objectives, the tasks are likely to be unrewarding for the children and mathematically impoverished. But if teachers plan from tasks, the activity is likely to be unfocussed and unassessable.

We propose two constructs to discuss task design. The first construct we call *purpose*. We define a purposeful task as one that has a meaningful outcome for the learner, in terms of an actual or virtual product, or the solution of an engaging problem. The second construct is *utility*. We have found that it is possible to plan for opportunities for learners to appreciate the *utility* of mathematical concepts and techniques in the sense that they learn how and why that idea is useful by applying it in a purposeful context. We claim this approach stands in contrast to the conventional emphasis on how to carry out a technique.

The difficulty in planning lies in linking purpose to utility in such a way that there is a high probability that the learner will stumble across the utility of the mathematical concept as they engage in the purposeful activity. I see this notion as critically important in the successful design of the Chance-Maker task. The children enjoyed the idea of working on the concrete and meaningful task of figuring out which gadgets were not working properly and then attempting to mend them. However, purpose is insufficient. The design of the Chance-Maker tools made it almost impossible for the children to avoid ideas that lie behind the two principal concepts: (i) the law of large numbers (through the repeat tool and the graphing facilities), and (ii) distribution (through the workings box).

No doubt there are many ways of engaging children in work on these two key ideas and the Chance-Maker case is just one such method. Nevertheless, perhaps the most critical implication that can be drawn for our pedagogy is not limited to probability but applies to task design in the general practice of teaching mathematics.

Testing Personal Conjectures

We know from the literature that children (and adults) have many idiosyncratic ways of thinking about chance situations. Some have been outlined above. How are we to regard these ways of thinking? My theoretical framework (Pratt & Noss, 2002), built as a synthesis of the work by Noss and Hoyles (1996) and diSessa (1993), asserts that old pieces of knowledge co-exist with newer pieces of knowledge, either in a connected way or perhaps isolated from each other. In this sense, misconceptions can either be useful platforms for further learning in the sense that they become connected to new knowledge, or they need to be shown to have relatively little explanatory power.

The role of feedback is crucial here. If children are to be in a position to refute long standing beliefs, or, as I would prefer to say, if they are to have less reliance upon those beliefs, they need feedback that gives them good evidence of the weakness of their current ideas. The children need to be able to test out their personal conjectures and evaluate them.

In the Chance-Maker study above, the children were able to test out the notion that the strength affected the result of throwing the gadgets. I refer to the strength bar as a *redundant* control in the sense that, mathematically the children do not need it. However, psychologically it is crucial that the children are able to test out their personal conjectures.

In the context of probability, this is even more important than elsewhere in mathematics. At this level, probability theory is essentially a model for describing certain types of phenomena. When we experience those phenomena, we make judgements about them, which rarely receive feedback that might cause us to reflect on whether our judgements were correct. When we play games, we attend to the excitement of the game; we do not usually reflect on our strategy (cf. the outcome approach). In any case, we are not usually in a position to try out the sort of long-term experiments that might give us helpful feedback.

Whatever tasks we design to help children to understand key objectives like the law of large numbers and distribution, those tasks must provide a mechanism for the children to appreciate the power of these ideas compared to their own intuitions.

Large Scale Experiments

An appreciation of the law of large numbers cannot be realised without the facility to carry out long-term experiments. It seems, from my research, that some ten-year-old children may well not have appreciated how randomness behaves in the long term. Their tendency, it seems, is to carry out a small number of trials of an experiment when given the choice. Why would you do otherwise unless you had good reason? Indeed children seem to follow a Law of Small Numbers (Kahneman et al., 1982).

My research suggests that the tasks that teachers give children should encourage them to decide for themselves how many trials to use. If they are simply told to use a large number of trials then how are they to realise the problems in using small numbers of trials, which they may well have done given a free choice? Even so, the task needs to encourage them to try increasingly high numbers of trials, and the available tools need to facilitate such large-scale experiments.

Systematic Variation of the Context

The situated cognition movement (see for example Lave, 1988, 1991) argues persuasively that knowledge is deeply contingent upon the setting. The children in my research did however eventually reuse their situated abstractions. Why do I think they did this? It seems that, once the children had seen the lack of explanatory power of their own ideas, they would reconsider recently learned knowledge. Furthermore, as those ideas proved reliable across different contexts, the ideas took on a higher priority (here I lean heavily on the notion of phenomenological primitives or p-prims: diSessa, 1988)[1] and were more easily cued as sense-making devices.

I believe that the children were eventually able to connect across the gadgets because there were huge structural similarities. In a sense the only difference between them was the outward appearance, which to the child was highly significant but to the mathematician is irrelevant. The underlying tools were identical in each case and, of course, the gadgets were wrapped up inside the same microworld.

Although a difficult challenge, new pedagogies that provide different contexts for the same mathematical idea and offer similar tools within each context may prove more effective in helping children to appreciate the wide applicability of the two key concepts, the law of large numbers and distribution.

5. FINAL REMARKS

The four pedagogic implications, purpose and utility, testing personal conjectures, large-scale experiments and systematic variation

[1] DiSessa's work provides a detailed model of conceptual change in which knowledge is seen as fragmented - at least in its initial stages. Small pieces of knowledge, p-prims, are abstracted directly from experience. One example of such a p-prim could be characterised as "I push – it moves". P-prims have priorities attached to them, which determine how likely any particular p-prim is to be cued, and this cueing priority is in turn modified according to how consistent and helpful the p-prim turned out to be in practice. Gradually, through "tuning towards expertise", p-prims may become connected to each other, forming what we might think of as concepts. Although situated abstractions are at a much higher grain level than p-prims, I have found his model useful as a way of thinking about the co-existence of different, possibly contradictory, situated abstractions and the process of tuning that might increase the likelihood of activation of normative abstractions.

of the context, have been abstracted from research, which depended fundamentally on technology. Critics of technology-based research in this domain refer to how children might not believe in the randomness of the computer, which is after all only pseudo-random anyway. I found that it was important that the children were able to persuade themselves that they could not predict or control the outcome from the computer, nor that they could find patterns in the results. Under these circumstances, children began to believe that the computer was indeed generating random results. The idea that the numbers from the computer are pseudo-random does not in itself worry me. From a modelling perspective, the use of a stochastic model to describe the results of a computer random generator are no different from using such a model to describe the results from a dice or any other physical random generator. It is though worth reflecting on the special nature of technology in relation to each of the four pedagogic guidelines listed above.

Purpose and utility are ideas that have arisen naturally from the *constructionist* ideas of Papert and others (Harel & Papert, 1991). Papert has argued that building concrete or virtual objects is a particularly appropriate way of learning. The mending task in the Chance-Maker study fits easily into this paradigm. I have found in many studies that building activity (and mending is a particular form of building) does lend itself to the difficult task of linking purpose to utility. Building seems to provide a concrete focus that lends the activity more meaningful. Yet, with careful design of the task, the designer is able to optimise the chance that the child will engage with planted mathematical concepts. The child is then able to construct utilities for those concepts through using them in the building process. Elsewhere (Ainley & Pratt, 2002), we have talked about other fruitful areas to explore in order to link purpose and utility. Thus, although computer environments used in a constructionist way facilitate the resolution of the planning paradox, they are not unique.

The testing of personal conjectures is an especially difficult aim to achieve without the use of technology. Hard-pressed teachers are unable to provide sufficient feedback to satisfy knowledge hungry children, given permission to explore. Without access to technology, teachers would have to employ other techniques such as group work and whole class sessions. The difficulty then is that the feedback is not neutral, it is less personal and it may be incorrect. On the other hand, computers are not very good at handling the range of possible idiosyncratic ideas that a child may hold, whereas people are far more flexible.

Large-scale experiments are also difficult to handle without technology. To avoid tedium, teachers tend to collate the work of whole classes. Unfortunately there is then some loss of individuality which means you may find yourself exploring someone else's way of thinking about the situation rather than your own. It seems to be a good strategy to begin with group work away from the computer in order to generate ideas related to conventional random generators. This process is about making explicit a range of personal conjectures. In order that these conjectures can be tested in a large-scale way, individual or small group exploration on the computer is likely to be needed as a follow-up to the initial group work.

Setting up different contexts might on the face of it look easier to do away from the computer, and there is some truth to this. However, it is important that the contexts have the same mathematical structure and probably some similar surface features to cue the reuse of recently learned knowledge.

One huge dilemma remains. When children use computers, will what they learn about probability "transfer" to conventional settings. Said in another way: "Will children reuse the knowledge in the new setting?" The above research suggests not, at least not in any simplistic way. If we think of a task set in the physical world as just another gadget, we see the problem. This gadget is so very different from those on the computer. The affordances of the physical world are in many respects insufficient for most people to gain rich intuitions for the key ideas. In the end, this is why the literature is full of reported failures. However, I hope that the suggestions above point towards the re-design of pedagogies that might bring virtual and the physical settings closer together, and enable children to reuse in conventional settings ideas initially constructed in the virtual world.

6. REFERENCES

Ainley, J., & Pratt, D. (2002). *Purpose and utility in pedagogic task design.* Paper presented at the 26th annual conference of the International Group for the Psychology of Mathematics, University of East Anglia, Norwich, UK.

Cobb, P., Confrey, J., diSessa, A., Lehrer, R., & Schauble, L. (2003). Design experiments in educational research. *Educational Researcher, 32*(1), 9-13.

DfEE. (1999). *The national numeracy strategy: Framework for teaching mathematics from reception to year 6*. Sudbury, Suffolk: DfEE Publication.

DfEE. (2001). *The national numeracy strategy: Framework for teaching mathematics years 7, 8 and 9*. Sudbury, Suffolk.: DfEE Publication.

diSessa, A. A. (1988). Knowledge in pieces. In G. Forman & P. Pufall (Eds.), *Constructivism in the computer age* (pp. 49-70). Hillsdale, New Jersey: Lawrence Erlbaum Associates, Inc.

diSessa, A. A. (1993). Towards an epistemology of physics. *Cognition and Instruction, 10*(2 & 3), 105-226.

Fischbein, E. (1975). *The intuitive sources of probabilistic thinking in children*. Dordrecht, The Netherlands: Reidel.

Fischbein, E., & Schnarch, D. (1996). *Intuitions and schemata in probabilistic thinking.* Paper presented at the 20th annual conference of the International Group for the Psychology of Mathematics Education, Valencia, Spain.

Harel, I., & Papert, S. (1991). *Constructionism*. Norwood, New Jersey: Ablex.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgement under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80*(4), 237-251.

Konold, C. (1989). Informal conceptions of probability. *Cognition and Instruction, 6*(1), 59-98.

Lave, J. (1988). *Cognition in practice*. Cambridge: Cambridge University Press.

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

Lecoutre, M. P. (1992). Cognitive models and problem spaces in “purely random” situations. *Educational Studies in Mathematics, 23*(6), 557-568.

Noss, R. (1997). *New cultures, new numeracies* (Inaugural Professorial Lecture Inaugural 14). London: Institute of Education, University of London.

Noss, R., & Hoyles, C. (1996). *Windows on mathematical meanings: Learning cultures and computers*. London: Kluwer Academic Publishers.

Piaget, J., & Inhelder, B. (1975). *The origin of the idea of chance in children* (L. J. Leake & P. D. Burrell & H. D. Fischbein, Trans.). London: Routledge & Kegan Paul. (Original work published 1951)

Pratt, D. (1998a). *The construction of meanings in and for a stochastic domain of abstraction.* Unpublished PhD thesis, University of London, London.

Pratt, D. (1998b). The co-ordination of meanings for randomness. *For the learning of mathematics, 18*(3), 2-11.

Pratt, D. (2000). Making sense of the total of two dice. *Journal for Research in Mathematics Education, 31*(5), 602-625.

Pratt, D., & Noss, R. (2002). The micro-evolution of mathematical knowledge: The case of randomness. *Journal of the Learning Sciences, 11*(4), 453-488.

Smith, J. P., diSessa, A. A., & Rochelle, J. (1993). Misconceptions reconceived - a constructivist analysis of knowledge in transition. *Journal of Learning Sciences, 3*(2), 115-163.

Tall, D., & Vinner, S. (1981). Concept image and concept definition in mathematics, with special reference to limits and continuity. *Educational Studies in Mathematics, 12*, 151-169.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5*, 207-232.

Wilensky, U. (1997). What is normal anyway? Therapy for epistemological anxiety. *Educational Studies in Mathematics, 33*, 171-202.

HOW DO TEACHERS FOSTER STUDENTS' UNDERSTANDING OF PROBABILITY? 201

Pratt, University of Warwick, United Kingdom