Designing for mathematical abstraction

Dave Pratt and Richard Noss

University of London, UK

Professor D. Pratt, Institute of Education, University of London, 20 Bedford Way, London, WC1H 0AL, UK. email: d.pratt@ioe.ac.uk (author to contact).

Professor R. Noss, London Knowledge Lab, Institute of Education, University of London, 23-29 Emerald St, London, WC1N 3QS, UK. email: r.noss@ioe.ac.uk

Abstract

Our focus is on the design of systems (pedagogical, technical, social) that encourage mathematical abstraction, a process we refer to as *designing for abstraction*. In this paper, we draw on detailed design experiments from our research on children's understanding about chance and distribution to re-present this work as a case study in designing for abstraction. Through the case study, we elaborate a number of design heuristics that we claim are also identifiable in the broader literature on designing for mathematical abstraction. Our previous work on the micro-evolution of mathematical knowledge indicated that new mathematical abstractions are routinely forged in activity with available tools and representations, coordinated with relatively naïve unstructured knowledge. In this paper, we identify the role of design in steering the micro-evolution of knowledge towards the focus of the designer's aspirations. A significant finding from the current analysis is the identification of a heuristic in designing for abstraction that requires the intentional blurring of the key mathematical concepts with the tools whose use might foster the construction of that abstraction.

It is commonly recognized that meaningful design constructs emerge from careful analysis of children's activity in relation to the designer's own framework for mathematical abstraction. The case study in this paper emphasizes the insufficiency of such a model for the relationship between epistemology and design. In fact, the case study characterises the dialectic relationship between epistemological analysis and design, in which the theoretical foundations of designing for abstraction and for the micro-evolution of mathematical knowledge can co-emerge.

Designing for Mathematical Abstraction

Our focus is on the design of systems (pedagogical, technical, social) that encourage mathematical abstraction, a process we refer to as *designing for abstraction*. We re-present previous work on young children's changing meanings for randomness as a case study in designing for abstraction. Through the case study, we elaborate a number of design heuristics that we claim have broader significance in the literature on designing for mathematical abstraction. In searching for descriptions and explanations that better fit children's activity in relation to our design activity, we treat theory building as a modelling exercise through which it appears that the conventional perception of the relationship between epistemology and design is insufficient.

The aims of the original research that led to this paper can be expressed quite simply. We aimed to research young children's changing meanings for randomness as those meanings evolved and were shaped by the digital tools that we made available to them. Since our concerns were to "address theoretical questions about the nature of learning in context" (Collins et al, 2004), we adopted a design research approach, culminating in a software tool, ChanceMaker.

In earlier work, we reflected on the children's meaning-making for randomness and distribution (Pratt, 2000), as it emerged through the design research, ultimately presenting a model for children's micro-evolution of knowledge (Pratt and Noss, 2002). What was inescapable during that research was that we not only learned about the children's meanings for randomness and distribution but we also became sensitized to the specific manner in which we could design tools that could act as a window on the children's meaning-making.

In this paper, therefore, we now seek to focus explicitly on the rationale for the specific design decisions taken in that research and to identify heuristics to guide future design activity.

We will present the case study chronologically since it is through the narrative that the relationship between epistemology and design becomes apparent. We begin by looking back at how we thought about epistemology and design at the start of the journey.

Initial orientation towards epistemology and design

When this research began (in 1994/5), the predominant view reported people's understanding of randomness and probability as essentially misconceived and that matters of chance were mostly counter-intuitive. The fallibility of people's conceptions of chance was reported, for example, in the seminal work by Kahneman and Tversky (for example, Kahneman, Slovic and Tversky, 1982), who catalogued the heuristics people used, normally in conditions where subjects were required to make swift judgements without recourse to tools or indeed teachers. Kahneman and Tversky demonstrated how, under such circumstances, the heuristics people used were subject to systematic bias, resulting in errors and evidence of misconceptions. The view that people's probabilistic comprehension is misconceived was reinforced by many other research reports, including those of Lecoutre (1992) and Konold (1989), who gave alternative explanations for the tendency of people to use a non-probabilistic approach (the *equiprobability bias* and the *outcome approach* respectively), in which such situations were regarded as mere matters of happenstance or luck.

Our own position was strongly influenced by the ideas of Smith, diSessa and Rochelle (1993). They asked if building new knowledge required replacing

4

misconceptions by better more normalised ways of thinking, what would be the foundations of such new knowledge? Replacing misconceptions unproblematically without any existing basis makes an anti-constructivist assumption and the search for misconceptions might support designers and teachers who need to be aware of existing conceptual structures that have developmental potential. We began our project with a belief that the misconceptions literature was reporting the effect of a particular research method and that by providing new tools, we might discover new ways to think about children's understanding of randomness. We recognised that our design research would need to sensitise us to what children might already understand about randomness so that we might build tools that could perturb that understanding and promote the development of more sophisticated knowledge.

In the literature on children's understanding of randomness and probability, there were notable exceptions to the predominant misconceptions perspective, in particular the critique in Wilensky's PhD thesis (1993) of the misconceptions approach and to some extent, the work on intuitions of chance by Fischbein (1975). Our orientation at the start of this research was consistent with that of Wilensky, aligned to the Constructionist school (Harel & Papert, 1991), which claims that, by offering microworlds in which children build meaningful products, the children will feel in control of the tools and become engaged in learning the mathematical ideas designed into the microworld. Constructionism suggests the efficacy of designing *quasi-concrete* virtual objects, instantiations of powerful mathematical ideas. By manipulating these on-screen objects, children can gain a sense of the power and limitations of the embodied mathematics (Papert, 1996), paralleling the construction of knowledge through the use of everyday artefacts in the material world:

"The principle is called the power principle or 'what comes first,

using it or 'getting it'?' The natural mode of acquiring most

knowledge is through use leading to progressively deepening

understanding. Only in school... is this order systematically inverted.

The power principle re-inverts the inversion."

(p. 98)

Although our focus was on research rather than teaching, we believed that full

engagement in exploratory activity would provide a window on children's

meanings for randomness, even as those meanings changed during activity (Noss

and Hoyles, 1996).

The emergence of ChanceMaker

In this section, we describe the chronological development of ChanceMaker up to

but not including its final stage. Even in our early attempts to bootstrap the design

process, we had in mind a particular setting in which the evolving designs would

be used.

The children who would use the tools were aged between 9.11 and 11.1

years in a primary school. By deciding to work with children from this school,

which was relatively advanced in its deployment of technological resources, the

distraction of technical obstacles would be minimised, allowing a tight focus on

the children's intuitions of the stochastic.

The initial bootstrapping idea to trigger the iterative design process in the

ChanceMaker study envisaged a simulation of a popular game, typically played at

school fetes; players roll a penny down a slope and win if the coin avoids landing

on any of the lines running across its path. The children were cast as the owners of

the roll-a-penny stall in the school fete. We hoped that such a task would be seen

as purposeful by the children in keeping with our Constructionist stance. The software design (see Figure 1) featured facilities to control the design of the sloped board (the `controls`), the entry fee and the prizes (the `money`).
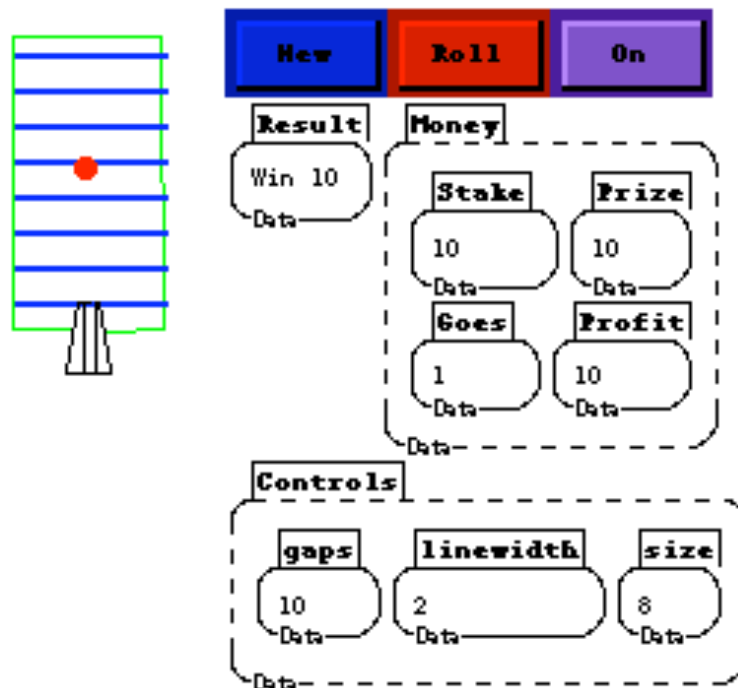

Figure 1: An early version of ChanceMaker based on a popular game played in school fetes. The prototypes of ChanceMaker were developed in Boxer, a Logo derivative that allows the tools to be re-programmable by the user but at the same time provides structures through which the designer can choose to make some aspects more immediately visible than others.

The player clicked the `roll` button and the animated coin rolled down the slope. As more coins were rolled, a distribution of results would be seen on the board as each attempt left a trace. Likely positions for the coin to land would show up as more dense patches on the board, not only emphasising variation in distances travelled but also indicating a distribution of distances through the density of the trace at different positions.

The children were intrigued by the task of designing a profitable game and were further motivated by the idea that they might build an actual game, based on

their prototype, for the coming school fete. As a window on the children's knowledge about randomness and distribution, the task was ineffective. Their focus was on the game itself and on financial calculations rather than on randomness and distribution.

The design was therefore modified to emphasize the idea of statistical distribution, by giving the children control over the movement of the coin (Figure 2). The `howfar` tool is a placeholder for the children's programming commands that, when executed, determine how far the coin will roll. Children were able to enter any legal code into the `howfar` tool.

Figure 2: The facility to control the movement of the coin was introduced

The intention was to encourage a focus on the relationship between the `howfar` box and the distances rolled by successive coins. However, the children were somewhat underwhelmed by the task! The motivation witnessed in the previous iteration had dissipated; there seemed now to be no driving force behind the children's activity. The children would carry out the task, responding politely to the researcher's prompts and exhortations, but there was no sense of urgency in them needing to find out and explore for themselves. In this respect, the task again was ineffective as a window on their thinking-in-change about randomness and distribution. It seemed that there needed to be a certain level of desire in order that the window might not simply reflect the researcher's agenda and in order to challenge the limits of the children's knowledge.

We continued to search for a design that would encourage children to rediscover the sense of engagement without losing the focus on randomness and distribution. We decided that the `howfar` box in Figure 2 was insufficiently expressive and we might be able to introduce a greater sense of engagement if the children had more avenues for expression. It was difficult to see what we should build into the `howfar` box by default. We used the primitive, `random`. If we had used another primitive it was difficult to see how the children would introduce randomness without us telling them to do so, and if we put the command `random` in there ourselves as default, we were leading their thinking too strongly and we would probably not gain a refined sense of their understanding of randomness.

We wondered what would happen if we asked the children to use their everyday knowledge of other random situations (or rather situations that seemed to involve randomness to us and might do so to the children). We built gadgets, simulations of everyday random generators, such as dice, lotteries, spinners. Figure 3 shows three of these new objects, now referred to as gadgets, alongside the `roll-a-penny` device.
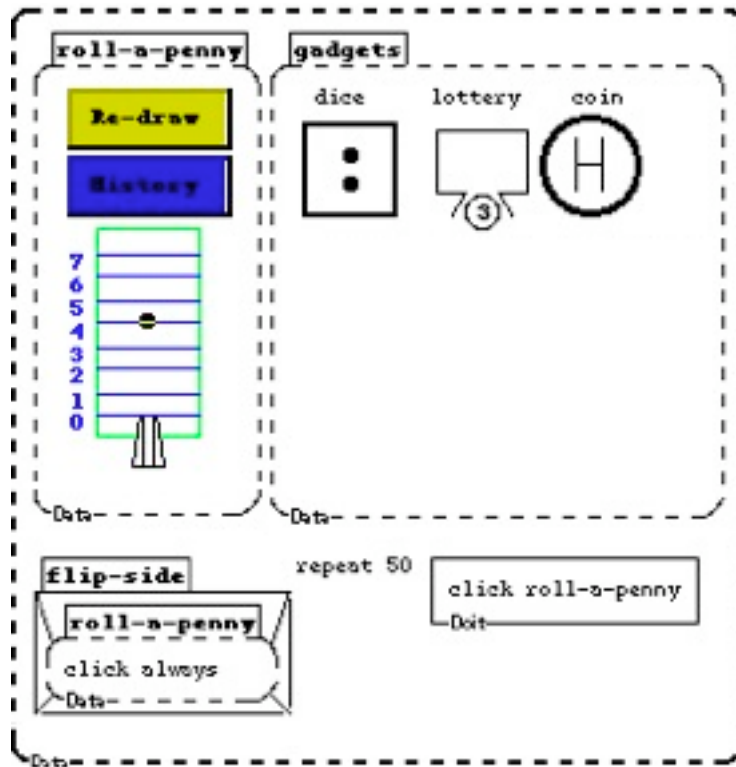
Figure 3: We introduced a range of devices, simulators of everyday experiences of randomness,

for example, tossing a die or a coin and drawing from a lottery. The children used these simulators

to express how they thought the roll-a-penny gadget should work.

In this design, the children were required to model the behaviour of the

roll-a-penny device using a gadget as the controller. Each gadget could be

activated by a command such as click dice or repeat 50 [click

dice]. The behaviour of any gadget could be changed by editing its flipside.

Figure 4 shows the flipside of the die gadget.

Figure 4: Each device could be flipped to reveal a flipside, which indicated how the device

worked. Here, the computer 'chooses' at random from 1 to 5. The flipside could be edited so the children might decide, for example, that the die should choose from 1 to 6.

By default, the `roll-a-penny` always went to the same distance of 4. We anticipated that the children would not believe that the `roll-a-penny` device should behave in this way and would search for a better model. For example, if the children believed that the `roll-a-penny` should behave somewhat like a die, they might edit the `roll-a-penny` flip-side to read `click dice` and then further edit the `dice` gadget to cover a range from 0 to 7 as appropriate for the `roll-a-penny` device. Essentially, the gadgets acted as a more expressive language to control the behaviour of the `roll-a-penny` device.

Although the expressiveness of the design was improved, when we tested this design, it seemed that the task was too obscure. Although we had not yet found a solution to the problem of designing for both engagement and focus, we had in fact stumbled upon what turned out to be a key development in the task design, which we will pick up again later.

Thoughts about epistemology and design, based on the iterative development of ChanceMaker

The experience of these early iterations led to a re-evaluation of our position on epistemology and design, which needed to accommodate the experience of observing children's levels of engagement and focus. Such factors clearly impacted on our ability to make inferences about meaning-making for randomness. In fact, these insights led to parallel research activity on a pair of connected constructs, *purpose* and *utility* (Ainley, Pratt and Hansen, 2006).

A *purposeful* task has a meaningful outcome for the learner in terms of an actual or virtual product, the solution of an engaging problem, or an argument or justification for a point of view. Designing tasks that engender such purpose is necessary but not sufficient for children to engage with the *utility* of ideas. In fact, the purpose of a task, as perceived by the learner, may be quite distinct from the intended learning outcomes. The purpose creates the necessity for the learner to use mathematical knowledge in pursuing the task and a benchmark for the child to recognise progress. Because the mathematical ideas are being used towards this end, children might come to appreciate *utility*: how and why the mathematics is useful, a form of understanding that has been largely ignored in the literature and is typically given little prominence in schools. Consider this example. Schools typically emphasize the learning of mode, median and mean as algorithms but an understanding of the utility of average would involve an appreciation of when average might be deployed as a useful idea. We regard gaining an understanding of *utility* as an essential element of mathematical abstraction.

Papert's Power Principle states that technology can allow children to learn mathematics through its use. We came to understand that, in order to build an appropriate window on the children's meanings for randomness, we needed to connect purpose and utility; not only must the task be seen as purposeful, as was the case with our first bootstrapping design, but also randomness and distribution must be seen as having utility. At this stage, we had not succeeded in connecting purpose and utility but the failure of the designs to date had opened up for us the potential significance of utility in children's abstracting process.

At about that time, we had also in fact stumbled upon a break-through idea: whereas previously we saw gadgets as mechanisms for expressing how the roll-

`a-penny` device should operate, the gadgets could themselves be broken and perhaps mending the gadgets could become the sought after purposeful task. The idea of mending a gadget took hold of our imaginations since it promised a constructionist "building-like" approach towards the study. We could reduce the obscurity of the previous iteration by removing the need to connect the other gadgets to the `roll-a-penny` device and simply offer all the gadgets as potentially broken. The challenge would then be to identify which gadgets were 'working properly' (that is to say, not broken) and to mend them. By careful design of the tools for mending, we intended to offer a purposeful task that would be likely to lead to the construction of utilities for randomness and distribution. The lack of clarity in the meaning of 'working properly' acted as a prompt for children to test out personal conjectures and provided a window through which the researcher might appreciate how the children thought coins, spinners and dice *should* behave.

When the children felt they had some appreciation of which gadgets were not working properly, they would be challenged to mend them. Indeed, as the task involved the children in making decisions about the performance of the gadget and about how to modify that performance, the task would we hoped, be seen as playful and engaging and encourage the children to be careless about their own naivety.

In summary, our initial epistemological position oriented us towards a constructionist perspective on the mathematical knowledge but the early iterations in the emergence of ChanceMaker enabled us to recognize the fundamental need to design a connection between purpose as might be construed by children and the utility of randomness and distribution. By stumbling upon the key role that mending gadgets might play, we understood that knowledge about randomness

was defined by knowing what was not random (for example, we might regard a process as random because we are unable to predict its outcome). We describe in more detail in the next section what this really meant about children's construction of meanings about randomness but even at this stage we understood that offering opportunities for mending gadgets would provide a window on how they anticipate randomness and so reveal their meanings for what was and what was not random.

The maturation of ChanceMaker

The task associated with the mature version of ChanceMaker was presented to the children in two parts. The first section below focuses on the initial challenge. Later, we describe the second part of the task. First then, it was explained that we were developing gadgets to behave like various familiar objects such as coins, spinners and dice. The children were asked to help us by playing with these gadgets and identifying which ones they thought were behaving 'properly'.

Identifying which gadgets were not working properly

In the mature version of ChanceMaker[1], we offered a range of gadgets; three are shown in Figure 5.

---

[1] The version shown here is coded in Imagine Logo (http://www.r-e-m.co.uk/logo/?comp=imagine)
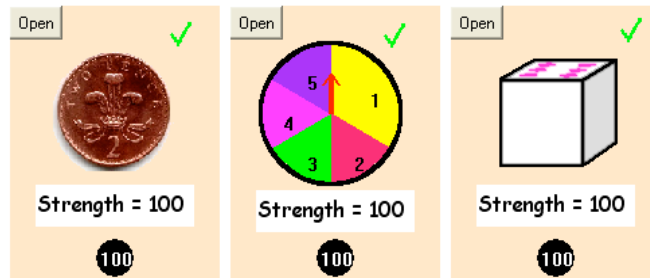
14

Figure 5: Three gadgets from ChanceMaker were the `coin`, the `spinner` and the `dice`. In each case, the strength tool could be pulled and let go in order to create an animated representation of 'throwing' the gadget.

The children were able to simulate the throwing or tossing of the gadget, either by pulling on the strength control, the black disc beneath the gadget itself, or by clicking directly on the gadget in order to replicate an experiment using the same strength as last time. The task was to decide which of the gadgets were working properly and which were not.

When children worked with ChanceMaker, we observed them articulating four expert-like meanings for randomness. In each instance, we give one example for illustrative purposes, though there were many such occurrences in the original data:

(i)   Unpredictability – when the children were unable to predict the next outcome, they would tend to regard the gadget as random. For example, in response to the question, "Do you think there is any number which is harder to get than any other number?" one child commented, "No . . . because it just comes out at random and any number could come out at any time so you don't really know which one is going to come out or which one is not going to come out."

(ii)   Unsteerability – when the children were unable to control the next outcome, they would typically view the gadget as random. For example, when asked to summarize how he would decide on randomness, one child responded, "Well, you decide by . . . if you're not controlling it or if you're not affecting it by doing anything, and if it's like not bad weather or anything or nothing's blowing it over or anything, that will be quite random, but if the wind was blowing it or you were putting force or it or something then it's not that random."

(iii)   Irregularity – when there appeared to be no sustainable pattern in results, the children would regard the gadget as random. For example, when asked how he would test the fairness of a dice, one child answered, "Testing it, I'd roll it and if it kept on going on one or another then I might think it's got like a magnet or something inside it . . . I'd test it about ten, fifteen times."

(iv)   Fairness – when the gadget appeared to be fair, the children would tend to regard the gadget as random. For example, one set of questions in the interview was designed to ascertain how the child thought about two spinners, one of which had uniform sectors and another, which had unequal size sectors. On the first uniform spinner, children often expressed concerns that the spinner may not be unsteerable but nevertheless recognized that there was no particular bias towards one number. In contrast, the same child would often regard the non-uniform spinner as non-random. One child commented, "No, because whoever made this, made the one

and the three bigger so you'll get the one and the three most of the time."

We say expert-like because at face value their meanings for randomness appear consistent with much expert opinion. However, we observed the children shifting quite rapidly between the four meanings, choosing whichever was triggered by seemingly superficial (from the mathematical point of view) aspects of the gadget or random generator. For example, a child referred to the unsteerability of a device and moments later referred instead to its unpredictability; the fact that situations that were not controlled were often not predictable apparently encouraged children to express these resources interchangeably.

At the same time, they would often fail to problematise potential conflicts. Thus, certain spinners might have been regarded as unfair (and so not random), because the sectors were unequal, and at the same time irregular, unsteerable or unpredictable (and so random), because they were not exactly able to determine the outcomes. In practice, such conflicts, which might have led to a new conceptualization of randomness, were not articulated; the children would categorize the spinner as random if they happened to pay attention to the unpredictability and not random if they happened to consider the non-uniform configuration of the spinner.

Thoughts about epistemology and design based on identifying broken gadgets

In trying to make sense of this data in which children quickly switched between meanings according to superficial (from our point of view) contextual stimuli, we regarded these meanings as, at best, weakly connected pieces of fragmented

knowledge. Perhaps even more significantly, the children would not articulate meanings for long-term randomness.

It is in the literature on physics education that we found the most persuasive theory of conceptual change that accounts for such fragmentation. diSessa (1993) has proposed that knowledge is, at least initially, made up of many small unstructured heterogeneous particles of knowledge, which he calls *phenomenological primitives*, or p-prims for short. P-prims are abstracted directly from experience. One characteristic of p-prims is that they contain priorities that shape how and when those pieces of knowledge are used to make sense of phenomena. As p-prims become recognized by the organism as more reliable in relation to their consistency, they are more likely to be called upon in the future as tools for abstracting. Gradually, through a process diSessa calls *tuning towards expertise* (1993)*,* clusters of connected p-prims will begin to appear simultaneously triggered by similar phenomena. diSessa refers to these clusters as *coordination classes*, roughly analogous to concepts. We were impressed by the extent to which this theory provides explanatory power for the inconsistency both between and within children (see, for example, Pratt & Noss, 2002).

In order to promote tuning towards expertise, we needed to allow and encourage the children to test the limits of their abstractions, that is, to test out their personal conjectures about the behaviour of the quasi-concrete objects that we were offering. Otherwise, how would children come to an appreciation of the lack of explanatory power in their current internal meanings about long-term aggregated behaviour? The aim then was not to classify such abstractions as misconceived but to support children in recognizing that they do not have powerful explanations, much in the manner of classical cognitive conflict approaches (Piaget, 1963).

We saw the second part of the task, mending the broken gadgets, as an opportunity to use new tools that would build in feedback mechanisms that were designed to increase the likelihood that children would identify the lack of explanatory power of their fragmented knowledge pieces in certain situations.
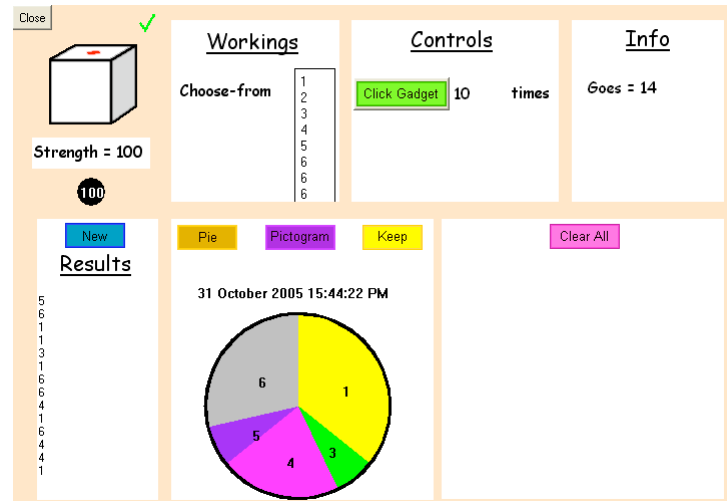

Figure **6:** The workings of the die gadget allow the child to continue playing at top level by clicking on the image of the gadget but also to use the tools to mend the gadget. The tools include a list or chart of previous results, the possibility of keeping a chart to compare later with a new chart, the total number of throws so far, an option to repeat quickly many throws (how many can be set by the child) and, most importantly, the facility to edit the workings of the box. Here, the die is choosing randomly with a bias towards throwing a 6. A child might edit this so that the die chooses from 1 to 6 with no bias.

Figure 6 shows the mending tools inside the die gadget. The other gadgets were similarly organized. The children were able to continue to play with the gadget itself as if at top-level or repeat quickly many throws of the die. It was possible to inspect the results or indeed to graph those results. In Figure 6, the child has simulated the roll of the die 10 times. The results can been seen as a list in the Results Box or as a graph, here depicted as a pie chart.

The `workings box` was introduced to the children as how the gadget works. So, in Figure 6, the die "chooses" between 1, 2, 3, 4, 5, 6, 6 and 6. No further detail on how the die might do this was offered and none was asked for. The researcher also demonstrated how the `workings box` could be edited so that the gadget could have a different set of choices. Thus, in Figure 7, the `workings box` has been edited so that the die chooses, in more conventional style, between 1, 2, 3, 4, 5 and 6. The child has simulated the throwing of the die 10 times and a pie chart has been generated. Typically, on seeing such a graph, a child might decide there are too many 6's in the `workings box`!
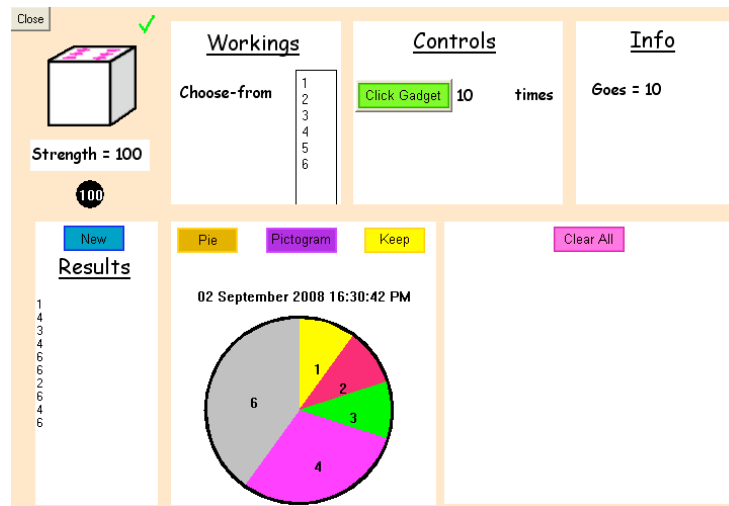

Figure 7: Here, the child has modified the workings of the die gadget so that the die chooses randomly from 1 to 6. However, the pie chart, after ten throws only, seems to indicate a bias towards 6, and it would be typical now if the child were to continue editing the working box, perhaps by removing the 6!

Thus, two important feedback mechanisms were i. the facility to observe the behaviour of animated gadgets and ii. to examine the lists of results and the graphs, such as pie charts and pictograms. Of course, children regularly ignore or do not perceive the significance of feedback but by setting the activity within a

20

purposeful task, we expected the desire to understand would minimize that possibility.

In making our design decisions, we sought to emphasize the lack of explanatory power in certain situations of the child's fragmented knowledge (in the lists below, the design decision relate to the final design shown in Figures 6 and 7):

- We made available prior results so that children might realize that apparent patterns were not sustained over a longer period.

- Children, who felt that the order of the numbers on the spinner was important, could edit the `workings box` to test whether different orders of numbers would influence the `spinner`'s behaviour.

- We expected from our early designs and from the literature that some children would predict that the outcome from the `coin` and many of the other gadgets to be influenced by how they threw the gadget. Hence, we provided a strength control through which they could throw the gadget with differing strengths. In fact, the strength mechanism was a *redundant control*. In other words, the strength control had no influence over the mathematics although it did extend or reduce the time of the animation. Thus, for example, the `coin` would spin longer or shorter depending on the strength of the throw but whether it would land as a head or a tail did not in any way depend on the value of the strength. In practice, children needed to explore the strength control for extended periods before they rejected the idea that the value of the strength had an impact on the outcome.

Cobb et al (2003) have recently argued that design research offers an approach in which theory cannot only be tested but can do work. The above list demonstrates some instances of how design could empower children to test conjectures. The knowledge-in-pieces perspective on the micro-evolution of knowledge makes an important prediction: The fragmented knowledge could, under appropriate conditions, evolve structures through tuning towards expertise. This prediction inspired us to move beyond the commonly espoused model of cognitive conflict by seeking to support recognition of the power of certain situated yet normative abstractions and we designed our system to mobilize current knowledge towards this end.

In our experience, there is usually some element of what the child understands that can be regarded as a root of expert knowledge. We need to elaborate what we mean by 'root'. In diSessa's terminology, fragmented knowledge manifests itself through articulations, often apparently inconsistent in nature, that the teacher or designer might note. Some of that inferred knowledge might not be inconsistent with expert knowledge even though it is likely to be different in scope. The child's knowledge might be more situated or even too generally applied; its power might not be recognised in the face of competing fragments of knowledge. For example, the ChanceMaker children had an appreciation of the unpredictability, irregularity and unsteerability of randomness. How might these meanings be used as starting points to support a process of tuning towards expertise, which might respond to the conflict generated by the earlier design decisions? What decisions might we make to increase the likelihood of triggering p-prims out of which expert knowledge might develop?

In order to support tuning towards expertise, we needed to identify and exploit what we recognized as the children's current knowledge. We have already

described, for example, how the ChanceMaker children had an appreciation of short-term randomness, articulated through local meanings such as fairness. In order to facilitate tuning towards expertise, we needed to design opportunities for building on that knowledge, as described below.

- We sought to mobilize (in the sense of tuning towards expertise) children's keen understanding of fairness, as attached to the appearance of a gadget, so that their notion of fairness was re-attached to the appearance of the pie chart and the `workings box`.

- We also sought to mobilize the children's appreciation of unpredictability by facilitating a recognition of its limitation to only short-term randomness, introducing through the design of the tools the possibility of predictability in relation to the aggregate over a longer term. We intended that tuning towards expertise might involve a move whereby local meanings for randomness and fairness would be mobilized to construct global meanings for predictability in the long term.

- Furthermore, we hoped that the repeat tool might enable the children to try out many cases to test out whether their idea worked in the longer term and to discover that what had been unpredictable in the short term seemed in some aggregated sense to be predictable in the long term.

Mending the gadgets

The second challenge to the children was to mend the broken gadgets using the tools found by opening up the gadget. We observed the children attempting to mend the gadgets, in particular by using the workings box. Details of their activity have been reported elsewhere (see by way of more detailed illustration the activity

23

of Anne and Rebecca in Pratt & Noss, 2002). Here, bearing in mind our focus on design, we give a précis of that activity.

The children tried out different configurations and inspected the Results Box and the `workings box` to understand the consequence of their mending actions. Typically, they would notice how the pie chart was not uniform, which conflicted, it seemed, with their expectation of fairness. The children would edit the `workings box` accordingly. So, if the pie chart contained too many 6's, they might remove a 6 from the `workings box`. Sometimes, by responding in this direct way to the data, the children would under-represent an outcome in the `workings box`. So, a die with `workings box` that read

`choosefrom [1 1 2 2 3 3 4 4 5 5 6 6]` might by chance generate a surplus of 2's. In response, the children might edit the `workings box` to read

`choosefrom [1 1 2 3 3 4 4 5 5 6 6]`. It typically took substantial experimentation before they experimented with a `workings box` that read

`choosefrom [1 2 3 4 5 6]`. Even then they might respond to, say, a large number of 1's by editing the `workings box` once more to become

`choosefrom [2 3 4 5 6]`.

Eventually, the children would usually realize that

`choosefrom [1 2 3 4 5 6]` would only mend the `dice` gadget if the `dice` were thrown many times. We characterized the way children expressed this construction with the phrase, "the more times you throw the `dice`, the more even is the pie chart". When challenged by the researchers with a `workings box` which did not have a uniform configuration, the children were typically once more puzzled and, after further experimentation, would articulate (through mouse-

clicks in the medium of ChanceMaker as well as natural language) the idea that "the more even is the `dice's workings box`, the more even is its pie chart" and in some cases "the more times you through the `dice`, the more even is the pie chart, provided the `workings box` is even".

The children moved from one gadget to another and we were struck by how, with each new gadget, they typically began the search for meaning from the beginning, not using the heuristics that they seemed to have constructed from the previous gadgets. In fact, since the internal structure was identical across the gadgets, we were surprised as observers at the fresh struggle for meaning with each gadget. However, we also noted that when all else failed, when they could find no other explanation for the new gadget's behaviour, children would often "remember" that "the more times you throw the `dice`, the more even is its pie chart" and "the more uniform is the `dice's workings box`, the fairer is its pie chart", conjecturing that perhaps this rule works also for the new gadget. Indeed, as the children became familiar with more gadgets, the time needed for them to become familiar with the new gadget and re-use the same heuristics for making sense of the gadget's behaviour reduced. Eventually, some children were able to anticipate outcomes. In one case near the end of two hours of activity, for example, the children were challenged by the researcher, "What I had really wanted was for the ones to have a very good chance and the twos to have only a fairly good chance." One of the children immediately edited the `workings box` to read `choose-from [1 1 1 2 2 3 4 5 6]`. In the context of the earlier struggles to make sense of the gadget's behaviour, this was strong evidence that the children were using the workings box as a predictor of behaviour.

Thoughts about epistemology and design based on mending the gadgets

Activity with the `workings box` led to a sense of the utility of distribution, or at least this situated version of the broader mathematical concept, in that distribution could be seen as a stochastic control over how a gadget works, both in the short term and in the long term. In retrospect, we see that we designed for purposeful activity that was optimized for the construction of situated utilities for distribution because distribution in the form of the `workings box` acted as a stochastic control over the long-term behaviour of the gadgets.

At the outset, there was no evidence to suggest the children could use the `workings box` to anticipate outcomes. Indeed, most of the early activity was characterized by surprise when the actual outcomes failed to match expectations. Only with experience and the construction of relatively reliable situated abstractions did the children begin to find that what happened matched their predictions. At this point, the `workings box` became something more than a way of changing what might happen; it became a way of predicting what would happen. In that sense the `workings box` could be thought of as a representation or embodiment of a piece of mathematics, a formalism for describing probability distribution.

From the design perspective, the `workings box` seems to embody a fusion between control and representation. Whereas the utility of the `workings box` as a control was that it could change how the gadget behaved, the utility of the `workings box` as a representation was that it could describe how the gadget would behave. This ambiguous role for the `workings box` appeared to provide the link between informal activity and formal representation, but the link was only

activated in practice because of the design success in connecting purpose and utility.

We referred to the first heuristic, "the more times you throw the `dice`, the more even is the pie chart", as N, and the second heuristic, "the more uniform is the `dice`'s `workings box`, the fairer is its pie chart", as D. We also referred to the co-ordination of these two as N.D. In fact N, D and N.D are examples of what Noss and Hoyles (1996) call *situated abstractions*, in which abstraction is conceived, not so much as pulling away from context, but as a process of constructing mathematical meanings by drawing context into abstraction, populating abstraction with objects and relationships of the setting. We began to see connections between situated abstractions and diSessa's knowledge-in-pieces. In fact, in using situated abstractions, rather than p-prims, as our analytical construct, we apply the knowledge-in-pieces metaphor at a somewhat larger grain size. Whereas p-prims are largely subconscious and so small in granularity that they can scarcely be represented in language (though diSessa tentatively attempts to do this by referring, for example, to the p-prim "I push it, it moves"), we prefer to think of abstraction at a level more commensurate with learners' conscious, expressible and therefore observable articulations.

The knowledge-in-pieces framework makes a second prediction: children's abstraction would be sensitive to context, since small changes in the situation would trigger quite different p-prims.
We felt that our epistemological position could now begin to articulate under what circumstances children might re-use pieces of knowledge in new settings, and how knowledge, which is becoming tuned towards expertise, be "transferred" across such settings?

It appeared to us that situated abstractions are likely to remain un-cued in new situations precisely because the relevant pieces of knowledge are not connected to that situation. We found it helpful, therefore, to think of a *contextual neighbourhood* surrounding pieces of knowledge. The neighbourhood captures the domain over which the idea has been encountered and found to be powerful by the child in explaining the on-screen behaviour. The short version of what we propose is that mathematical abstraction is a process of broadening that contextual neighbourhood, rather than of de-contextualization.

Discussion and summary

We set out to elaborate a case study of design research, seeking to identify design heuristics that emerged. We also aimed to convey the dialectic relationship between designing for abstraction and the designers' perspective on the micro-evolution of mathematical knowledge. We now summarise the situation.

First we discuss the design heuristics that have been illustrated through this case study. A starting point was that a knowledge-in-pieces framework for understanding conceptual change led us to envisage the design task, not so much as supporting the replacement of misconceptions with normalized knowledge, as seeking to re-prioritise underlying p-prims.

We therefore designed to enable children to test out their personal conjectures and to support enhancement of priorities that might be attached to pieces of knowledge that we could envisage might lie at the root of normalized knowledge, i.e. standard knowledge, socially accepted by the mathematical community. Our window on children's knowledge had to be sensitive to the children's pre-existing meanings for randomness and distribution. In practice, we recognized the importance of *fairness*, a meaning for randomness articulated by

all of our sample children and generally associated with the immediate physical appearance of the gadgets. So we designed tools such as the pie chart to encourage realignment of the notion of fairness with that of equally likely outcomes.

To summarise, we can state the above design activity in terms of two emergent design heuristics:

1. Enable the testing by children of their personal conjectures.

2. Seek to enhance the explanatory power of knowledge that might offer a route to normalized knowledge.

In order to execute the second heuristic, we needed to design a task which would be likely to engage the child in activity, focussed on the key mathematical concepts of randomness and distribution. In effect, we designed a window on that mathematical activity through a task seen as purposeful by the child. The task was carefully designed so that in pursuing his aims, the child would be likely to come to appreciate the *utility* of distribution. Thus, the task design connected purpose and utility by inventing a novel representation for distribution and presenting it as a means of control over core mathematical activity. The children used each `workings box` as a control over the behaviour of the corresponding gadget; subsequently each `workings box` became a means for the children to predict the behaviour of the gadgets. In this sense, the `workings box` fused control with representation, enabling a blurring of the informal and the formal and providing a natural connection between the purposeful activity of the task and the key mathematical concept that we wanted the children to address. In short, a third and fourth design heuristic emerged:

3. Construct a task design that will be likely to generate purposeful activity and tools that encourage the construction of utilities for the key mathematical concepts.

4. Identify or design representations of key mathematical concepts that can be used as control points needed by the child to pursue their aim.

We can now revisit other work and try to identify post-hoc whether these four heuristics appear relevant to their designs. To explain what we mean, we turn to a familiar example for readers of this journal. Young children using the Logo turtle to draw pictures or create animations may already have fuzzy ways of thinking about distance or angle. By using commands such as `FD 50` and `RT 90`, they test out those personal conjectures and begin to revise their ideas through use. These commands are acting as controls over the drawing/animation activity and become representations in the sense that gradually many of these children will be able to predict the result of issuing the command and even create procedures where the result of the procedure is imagined in advance of execution. The utility of angle will be seen in terms of its power to produce pictures and more generally create or describe movement. What the turtle graphics commands like `FD 50` manage to do is to combine computation – it will do something potentially useful when run - with representation – what it will do can be predicted from the representation. In earlier work, we named this duality 'auto-expressive': "environments in which the only way to manipulate and reconstruct objects is to express explicitly the relationships between them" (Noss, Healy & Hoyles, 1997, p.5). The same is true of the `workings box`.

Put this way, it is now possible to understand the failure of earlier iterations of ChanceMaker, in which the gadgets were used to express how the `roll-a-`

`penny` device should behave. The immature gadgets were not representing some generic mathematical idea; rather they were limited to the functioning of an everyday phenomenon such as a coin or die. In this sense, they did not possess the expressiveness and generality of a mathematical representation. In contrast, the `workings box` in each mature gadget contained within itself the dual characteristics of controlling behaviour of the gadget and representing a powerful mathematical idea.

Our second aim for the case study was to articulate the relationship between epistemology and design. Our case study highlights the evolution of this reciprocal relationship. It is commonly understood that "the design of a tool reflects the designer's intentions, thoughts and compromises" (Yerushalmy, 1999, p. 171) but in our emphasis, in effect, we add the reciprocal that *the designer's epistemological understanding reflects the design of the tool*. Taken together, these two statements capture the dialectical nature of the relationship between epistemology and design. Below we trace the manner in which as a result of the design process our understanding of epistemology changed.

Our original stance, essentially Constructionist, became more refined through several stages. Our early experiences in trying to bootstrap the design research helped us to recognize the significance and non-triviality of designing a task through which the child's purpose for the task leads to the construction of utility for a key mathematical concept. Such insights have led us to review our stance on knowledge construction to incorporate the notions of purpose and utility as discussed in the section on *Thoughts about epistemology and design, based on the iterative development of ChanceMaker*. Designing for abstraction then has to

consider what children might find purposeful and to invent tasks that optimize the chance that such a *purpose* might lead to the construction of utility.

A second change in our epistemological understanding took place. We found that adopting diSessa's framework gave us insight into the evolution of the children's thinking in terms of the situated abstractions, N, D and N.D. Situated abstractions (SA) seemed to offer a way of thinking about knowledge at a grain size more aligned to the children's articulations than the sub-conscious behaviour captured by diSessa's p-prims. In contrast, but complementary to, the p-prim idea, the notion of situated abstraction privileges the explicit and expressed rather than the implicit and intuitive. In characterising expression-with-tool, the evolution of the children's thinking that led to the N, D and N.D. SAs were firmly embedded within the toolset. This naturally led us to asking what, if anything, might be 'transferred' in this process. While the evolution of conceptual change at the p-prim level involves reorganisation and evolution, we found it helpful in the SA analysis to think in terms of contextual neighbourhoods over which SAs were seen to apply. Of course the determination of contextual neighbourhoods was, as best we can judge, entirely implicit - even subconscious - much as diSessa's p-prims.

Thought of in this light, therefore, *designing for mathematical abstraction* becomes a challenge to create a domain for the articulation of situated abstractions, and the means by which the contextual neighbourhoods can be refined and expanded. The breadth or limitation of a contextual neighbourhood is appreciated through recognising (possibly sub-consciously) the relevance of prior knowledge to ongoing activity. By fusing control over mathematical activity with the representation of that mathematics, the designer offers the opportunity for

activation and transformation of the contextual neighbourhood as additional or restricted utility of the mathematical concept.

We have shown through this case study how designing for abstraction can generate new theory not only about design in the form of heuristics, as summarised in the first part of this discussion, but also, as in the second part, about the mutually generative nature of epistemology and design.

References

Ainley, J., Pratt, D. & Hansen, A. (2006). Connecting engagement and focus in pedagogic task design, *British Educational Research Journal*. 32(1), 21-36.

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

Collins, A., Joseph, D. & Bielaczyc, K. (2004). Design Research: Theoretical and methodological issues, *Journal of Learning Sciences*, 13(1), 15-42.

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*: Reidel.

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.

Konold, C. (1989). Informal Conceptions of Probability. *Cognition and Instruction, 6*, 59-98.

Lecoutre, M.P. (1992). Cognitive Models and Problem Spaces in "Purely Random" Situations. *Educational Studies in Mathematics*, 23, 589-593.

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

Noss, R., Healy L. & Hoyles, C. (1997) The Construction of Mathematical Meanings: Connecting the Visual with the Symbolic. *Educational Studies in Mathematics* 33: 203–233

Papert, S. (1996). An exploration in the space of mathematics educations. *International Journal of Computers for Mathematical Learning,* 1(1), 95-123.

Piaget, J. (1963). *The Psychology of Intelligence.* New York: Routledge

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.

Wilensky, U. (1993). *Connected Mathematics - Building Concrete Relationships with Mathematical Knowledge*. PhD Thesis, Massachusetts Institute of Technology.

Yerushalmy, M. (1999). Making exploration visible: on software design and school algebra curriculum. *International Journal of Computers for Mathematical Learning*, 4(2-3), 169-189.