

Running head: THE MICRO-EVOLUTION OF MATHEMATICAL KNOWLEDGE

The Micro-Evolution of Mathematical Knowledge: The Case of Randomness

Dave Pratt and Richard Noss

University of Warwick, UK, and University of London, UK

Author Note

Dr D. Pratt, Mathematics Education Research Centre, Institute of Education, University of Warwick, Coventry CV4 7AL, UK. email: dave.pratt@warwick.ac.uk. (author to contact)

Professor R. Noss, Mathematical Sciences Group, Institute of Education, University of London, 20 Bedford Way, London, WC1H 0AL, UK. email: rnoss@ioe.ac.uk

Abstract

In this paper we explore the growth of mathematical knowledge and in particular, seek to clarify the relationship between abstraction and context. Our method is to gain a deeper appreciation of the process by which mathematical abstraction is achieved and the nature of abstraction itself, by connecting our analysis at the level of observation with a corresponding theoretical analysis at an appropriate grain size. In this paper we build on previous work to take a further step towards constructing a viable model of the micro-evolution of mathematical knowledge in context.

The theoretical model elaborated here is grounded in data drawn from a study of 10-11 year olds' construction of meanings for randomness in the context of a carefully designed computational microworld, whose central feature was the visibility of its mechanisms--how the random behavior of objects actually "worked". In this paper, we illustrate the theory by reference to a single case study chosen to illuminate the relationship between the situation (including, crucially, its tools and tasks) and the emergence of new knowledge. Our explanation will employ the notion of situated abstraction as an explanatory device that attempts to synthesize existing micro- and macro-level descriptions of knowledge construction. One implication will be that the apparent dichotomy between mathematical knowledge as de-contextualized or highly situated can be usefully resolved as affording different perspectives on a broadening of contextual neighborhood over which a network of knowledge elements applies.

The Micro-Evolution of Mathematical Knowledge: The Case of Randomness

Mathematical meanings are elusive. We know surprisingly little about their genesis. It has become fashionable, almost obligatory, to display constructivist credentials, and acknowledge that mathematical meanings are built on existing knowledge, that children (and adults) do not come to learning situations as *tabula rasa*, and that meanings are constructed by the individual (or, depending on taste, the relevant cultural group).

The trouble is that this begs as many questions as it answers. It fails to address at a micro-level, the problem of how--if meanings are layered one on another--anything new is learned at all. What kind of grain-size do mathematical resources have? Are they like mini-theories of the way the mathematical world works? Or are they atoms of fragmentary mathematical meanings, essentially standalone? If the former, how is change possible? If the latter, what structures these atoms and how does restructuring occur? By asking these sorts of questions, we create a hostage to fortune, as we cannot pretend to do more than point to some partial answers to a subset of these questions, not least because they pose problems which are independent of the mathematical knowledge domain, while we focus on just one. In this paper, therefore, we position ourselves firmly in the domain of theory construction, building on previous theoretical work (particularly in Noss and Hoyles, 1996) and we will illustrate our case by reference to a single case study of two children, drawn from a considerably larger corpus of data (Pratt, 1998a).

Underpinning much discussion concerning the origin of mathematical knowledge lies the question of mathematical abstraction. It is a thorny question, whose problematic status has recently been addressed by a number of authors whose contributions we will consider below. A central issue is the extent to which mathematical abstraction depends on de-contextualization, or more generally, to explicate the relationship between abstraction and context. One of the difficulties surrounding the clarification of the relationship between

mathematical abstraction and de-contextualization, is that the finished “product”-- mathematical abstraction considered as an achieved state of thought--certainly appears to involve a state of having cut loose from contextual boundaries, and a breadth of application across settings. But this does not, in itself, deal sufficiently with the process. If mathematical abstraction means de-contextualization, does this involve a blurring of boundaries between contexts? Or a new cognitive mechanism to map abstraction onto different contexts? What, in any case, is it that is abstracted? Does mathematical abstraction involve a process of transfer, conceived in traditional psychological terms, in which knowledge in a single context is mapped out of old contexts and into new ones?

A common approach is to regard mathematical abstraction as a strictly hierarchical process, progressing through a series of stages. For example, Dubinsky (1991) forwards the notion of a “process,” an interiorized version of an action in which a repeatable physical or mental transformation of an object or objects takes place. At some point, a process can be transformed by an object, when, according to this theory, the process has been encapsulated to become an “object.” Sfard (see, for example, 1991) also regards the isolation of mathematical objects as a key achievement of mathematical abstraction.

These process/object models have taken steps beyond identifying and labeling stages of mathematical abstraction: For example, Sfard (1994), in a paper in which she adopts a richer and more sociolinguistic stance than that of her 1991 work, has described reification as the birth of metaphor, which renders the knowledge more integrative and manipulable; Gray and Tall (1994) refer to the ambiguous nature of process and object.

While these and other contributions represent a helpful development in the theoretical literature, they are nonetheless somewhat limited in their capacity to describe the process by which the various states of abstraction are achieved (other workers have similar misgivings: see, for example, Confrey and Costa, 1996). We believe that a further step would result from

reappraising the relationship between mathematical abstraction and de-contextualization and elaborating the process by which knowledge becomes abstracted. This project would need to pay particular attention to the role of the setting in the construction of mathematical knowledge.

One way of thinking about the role of setting has been strongly influenced by the paradigm of situated cognition. The strength of this view derives, in no small part, from the work of Jean Lave (1988) who showed convincingly how settings shape strategies for individuals, how problems are solved in the course of action, structured by activity. Overarching strategies or generalized solutions have little place in everyday experience; mathematics in situ has no need for universal laws, consistency, generality--it is concerned with getting the job done using whatever affordances are available. People do just enough to cope with the demands of the setting in which they find themselves. The strong case is made by Resnick, who argues that: “every cognitive act must be viewed as a specific response to a specific set of circumstances. Only by understanding the circumstances and the participants’ construal of the situation can a valid interpretation of the cognitive activity be made” (Resnick, 1991, p. 2).

We have argued elsewhere (Noss & Hoyles, 1996) that this view leads into a mathematical cul-de-sac. For if it is really true that every act is nothing more than a specific response to a specific set of circumstances, how does mathematical knowledge grow? What can be salvaged from the idea of mathematical abstraction? Mathematics is the science of structure and pattern, the study of relationships between relationships; and to that extent at least, it cannot be immutably tied to real referents, situations whose meaning is contingent upon the particularities of setting: at least not in the traditional sense of “real”.

Other workers have proposed alternative ways of conceptualizing mathematical abstraction. For example, Wilensky (1991) proposes that one should view the concreteness of

a concept as a function of the relationship between the learner and the concept. He suggests that, as one becomes more and more familiar with an idea, one makes more and more connections with it; one concretizes the concept. According to this view, ideas become increasingly concrete. Advanced mathematical concepts are neither more nor less abstract per se. For an individual who has not had the opportunity or does not yet possess the internal connections, a concept will be abstract, disconnected. That same concept may be quite concrete for another individual.

Noss & Hoyles (1996) go beyond considering the individual in the process of abstracting knowledge. Building on Wilensky's theme of connections, they extend the idea to considering a whole network of such links, encompassing not only the individual but also resources external to that person. We will use the word "resources" from now on, rather than variations on "meanings," "knowledge," "concepts," etc., to emphasize the complementary roles played by internal (cognitive) and external (physical or virtual) sources of meaning making. Thus the term resource will encompass both external tools and internal knowledge, including both informal knowledge like intuitions and formal conceptual knowledge. To describe this network, Noss and Hoyles use the idea of webbing to evoke the ways that learners come to construct new mathematical knowledge by forging and reforging internal connections through the interaction of internal and external resources during activity and in reflection upon it. The notion of webbing aims, therefore, to recognize the central significance of tools as external resources that shape the nature of the mathematical resources constructed, resources that have been observed to be highly dependent upon the particular attributes of those tools as cognized by students.

As a way of pinpointing what kind of knowledge emerges in the course of activity, Noss and Hoyles propose the notion of situated abstraction (see Hoyles and Noss, 1992; Noss and Hoyles, 1996). Situated abstractions emerge during activity as internal resources that

serve as relatively general devices for making sense of situations that arise within a setting. The idea tries to describe the types of knowledge that enable learners to reflect on the structures within a setting and make sense of phenomena that hold true across it. At the same time, a situated abstraction, or, to put it more exactly, the relationships and actions based upon the situated abstraction, are expressed in a language (not necessarily verbally articulated) that remains embedded in the situation in which it was constructed, potentially constraining its validity in new contexts, with different tools and affordances.

Formal abstraction is, by definition, characterized by its separation from the situation(s) that give rise to it. A crucial property of formal abstraction is that it occurs within a discourse which has its own syntax and semantics, and which is distinct from its genetic context. It is this which gives formal abstraction its appearance of containing objective essence, and which is responsible for the assumption that this essence exists in a form which allows it to be “transferred” to new situations (see Noss, Hoyles, & Pozzi, in press, for a study which uses the idea of situated abstraction in the context of understanding professional practices).

It is precisely that mathematical abstraction is generally expressed in a formalism with its own rules of transformation that affords it its extraordinary power. The closure of the formal system is crucial for this capacity: it arises because relationships within the formal system can be articulated without the “noise” that is characteristic of real systems. Situated abstraction, on the other hand, emphasizes connection with situations, not seeking to challenge the utility of formal abstraction, but maintaining that abstraction can take place in situ rather than only within a self-contained system mentally and/or physically remote from it.

The central premise is that individuals can and do find invariances which span settings, but this is not always best described as a mapping from a situation to some other mental realm and back again. Individuals can abstract knowledge within settings and remain tied to the objects and relationships within the situation, its tools, linguistic conventions and

structures. Situated abstraction is observable as more or less tacitly articulated invariance of relationships, framed within the situation itself.

By blurring external and internal resources, Noss and Hoyles bring into close proximity mathematical abstraction and its setting. Other workers have made similar observations, notably Nemirovsky (in press) who has described situated generalizing, as a recognition of the generalization process where there is no fixed boundary between experienced phenomena and “the realm where the generalization takes place.” His approach focuses a little more closely on the process of generalization, whereas, in this paper at least, our interest centers more on the conceptual domain.

The foregoing ideas offer variations on a theme that blur any strong distinctions between experience out there and mathematical abstraction in here. Whilst such theories provide constructs that elaborate the relationship between mathematical abstraction and de-contextualization at the macro-level, there is much to do in elaborating what processes of abstraction in situ might mean at a greater level of detail. Thus, in the remainder of this paper, our goal will be to refine macro-level descriptions of the emergence of new mathematical knowledge by studying the relationship between tool-use and situation, and in so doing, elaborate some of the micro-level details of the webbing idea.

In order to achieve this objective, we need to integrate the macro-level ideas of webbing and abstraction into existing micro-level descriptions of knowledge construction. This is what we mean by the term “micro-evolution” in the title. A good place to start is to ask how intuitive or primitive knowledge gets there in the first place. diSessa (1983) proposes that we abstract primitive intuitions directly from experience, that knowing is essentially phenomenological. He argues that “p-prims” (phenomenological primitives) are a collection of heterarchical and rich ways of seeing and sometimes explaining the world (e.g., things move in the direction you push them). They are relatively minimal abstractions of

simple common phenomena and are stored in what he calls the “causal net,” a network of primitive sense-making causal relationships. As learning develops, these p-prims become increasingly structured and “co-ordination classes” are formed. This structuring can be thought of as connections between p-prims that enable clusters of p-prims to be simultaneously (more or less) triggered by observed characteristics of phenomena. Such clusters roughly parallel what other writers might refer to as a “concept” but which diSessa calls a “co-ordination class” (diSessa, 1993).

According to diSessa, conceptual change occurs in three ways, each involving a transformation of the causal net: (i) by the addition of new p-prims, (ii) by the formation of new connections between p-prims, and (iii) by changes to the priorities which fix how likely a p-prim is to be triggered by incoming data (its cueing or reliability priority), in terms of its consistency with other p-prims or with further incoming data. The unstructured p-prims live in relative isolation, and although they may be triggered individually by phenomena, they will not then reliably trigger further p-prims. P-prims would remain unstructured if incoming data were not inconsistent with those p-prims or if there were no further incoming data. When p-prims are seen to be consistent with each other on a regular basis, they become highly structured so that a whole cluster is always triggered at the same time.

diSessa’s detailed model of conceptual change is helpful, not only because it offers clarity at both micro and macro-levels, but also because it makes provision, through the clustering of p-prims and through the mechanism of priorities attached to p-prims, for the construction, and possible evolution, of situationally-dependent knowledge.

While diSessa’s formulation is an attractive starting point for our own work, there are a number of issues that caution against a too-casual application of the theory to the task at hand. The first concerns diSessa’s focus on physics, which leads him to consider causality as the phenomenological basis for p-prims. This seems at first sight to stand in contrast with

mathematical knowledge, which seeks to express meanings through inference, independent of context. In what follows, we will not elaborate the causal/inferential distinction. More generally, we will have to be cautious about an oversimplified mapping of p-prims to situated abstraction: As observables at the macro-level, situated abstractions are relatively (to p-prims at the micro-level) complex and likely to depend upon collections of p-prims (and much more besides).

The second issue relates to the role of tools in shaping the construction of mathematical knowledge. It would be useful to elaborate the relationship between the cueing of p-prims (or any similar mechanism) and the perceived functions of tools within a setting. If we could understand this issue more clearly, we might be able to elaborate how the web is shaped to make some resources more likely to be used than others in any particular situation.

The third issue involves the problem of de-contextualization. DiSessa uses the term “tuning towards expertise” to describe the move towards more expert forms of knowledge. In mathematics, one characteristic of expert knowledge is its apparent cutting away from context. Any model of mathematical meaning-making should account for some process by which elements of knowledge take on more coherent structures, in which, perhaps, learners constructing situated abstractions come to coordinate them with each other, or with existing knowledge. We would like to address, therefore, how re-structuring might provide some insight into the decontextualization problem.

These three issues will guide our discussion in what follows, and we now consider each in turn.

The Problematic Nature of Context and Mathematical Abstraction

(i) What is abstracted?

The urge to describe how things work is, it seems, primitive to humans. The sense that contact with moving things conveys motion, or that pushing harder makes objects move

faster, relates to a recognition that the world works in a (mostly) causal way, that phenomena have explanations for the way they work. What is the equivalent insight for mathematics? Is the insight that “multiplication makes bigger” or--more relevantly for our own study-- “random sequences are patternless” causal in nature? Perhaps so. The distinction between inference and causality is, after all, an adult rather than a child distinction, and there is little reason to doubt that children’s beliefs about randomness are shaped by their experiences with throwing dice and tossing coins. P-prims, whether causal or inferential, are abstractions of reality, not mere descriptions of it. On the other hand, the intuition that, say, randomly-generated events give rise to patternless sequences does appear--at this stage we are being necessarily vague--to necessitate a degree of inference that, say, “force causes movement,” does not. We note, however, that diSessa (1993) recognizes that the genesis of p-prims “might be more complex than a simple abstraction schematization of a phenomenon” (p. 151), a statement which identifies the possibility that sense-making devices such as p-prims can be made out of other, more primitive devices.

We ought at this point to discuss the question of language in the process of knowledge construction and application, as we have several times in the foregoing discussion, loosely used terms like “articulation” and “expression.” In this respect, we need to consider the micro-level and macro-levels separately. From a micro-level perspective, it is reasonable to agree with diSessa that p-prims are not strongly associated with dictionary lexicon:

Much less do they have explicit propositional form. . . . Subjects may make predictions on the basis of a p-prim, but the prediction is not the p-prim.

(*Ibid.* p. 119.)

However, if we consider a larger grain size of knowledge, there may be some interesting differences. A situated abstraction is an articulated expression of invariant relationships perceived within a setting, and expressed within the tools and linguistic forms

of the setting. Note that this may, or may not, have much to do with dictionary lexicons: the linguistic structures may well have as much to do with mouse-clicks or gesture as spoken language. In this case, while we may certainly agree that the prediction is not the p-prim--or any other knowledge structure--it might sometimes be quite difficult to distinguish between the two. We will have more to say on this when we have illustrated our thinking with data on children's ideas about randomness below.

(ii) How is it abstracted? What is the role of tools?

We take it as axiomatic that the external resources available for making sense shape the knowledge development in fundamental ways. Our particular interest centers on the relationship between external tools and the set of internal resources--however structured--that the learner possesses at any time. If we wish to understand how new knowledge is developed, it might help to get a clearer picture of how specific pieces of knowledge and the relationships between them are activated by the experience of encountering the (real or virtual) world. This is, of course, a tall order. But a model that shed light on the relationship between external resources and the evolution of mathematical knowledge would have pedagogic value: We could gain a handle on how to provide tools and activity structures with optimal opportunities for supporting the construction of new knowledge.

This is, in essence, the fundamental design problem of mathematics education. And it is here that we believe that digital technologies hold particular promise. Our use of digital technologies in the study we report here is not at all centered on any a priori belief that we can build the right tool to "enhance understanding." Rather, our aim is to harness what Seymour Papert (1982) has called the "Protean" quality of digital technologies, the possibilities they afford for building and rebuilding virtual structures and to observe the ways in which people interact with them. This is what we mean when we call the computer a "window" on mathematical meanings (Noss and Hoyles, 1996): it is a screen on which we

and our students can express aspirations and ideas, and which — because these ideas must be made more or less explicit — afford considerable insight to the researcher in the process. It is in this sense that we claim the process of design as a contribution to the process of theory building.

(iii) Is mathematical abstraction de-contextualized?

We stated earlier that de-contextualization appears to play a crucial role in the process of mathematical abstraction. Yet all of the literature on situated cognition and much besides, shows just how problematic this is, as even without the strong version of the situated cognition thesis, we accept that knowledge in general tends to remain tied to the situation that gave rise to it. If we consider that phenomenological knowledge exists as some kind of causal network, then it appears incontrovertible that it maintains a strong link to the original primitive intuitions abstracted directly from experience.

From the point of view of mathematics, two questions arise. First, what exactly is de-contextualization? It will be necessary to refine what we mean by a mathematical context, and try to characterize what we mean by situational specificity in order to see whether and how learners sometimes break out of contexts (and, of course, whether “break out” is the right metaphor at all).

A second question concerns abstraction. The etymology of the word speaks for itself: the assumption is that abstraction demands a pulling away from context. Once achieved, there is no doubt that from a mathematical point of view, it is this that lends mathematics its power as a means to describe the world. But how is it achieved? In what circumstances? And, more importantly, is it a binary state system, in which one either has or has not abstracted? Whatever the case, we ought to be able to say something about mathematical abstraction-in-progress, particularly if we set things up so that we can study learning in a relatively transparent way within appropriate computational settings.

We now wish to set out as clearly as possible our goal, which in general terms is to develop a model of the micro-evolution of mathematical knowledge in context.

Elements of a Model for Mathematical Abstraction in Context

The study of abstraction must involve the consideration of currently accessible external and internal resources. Even if, finally, mathematical abstraction appears to be de-contextualized, it would be remiss to assume that context plays no part in the process of abstracting. Indeed, there is every reason to believe the role of tools is fundamental, begging the question as to how mathematical abstraction comes to appear cut off from context. This approach to the study of mathematical abstraction leads us to the conclusion that a model for how mathematical knowledge grows must contain five essential elements:

1. A clear description of the nature of naïve mathematical knowledge including the grain size of emergent mathematical knowledge.
2. A description of the setting, including structuring resources designed to perturb thinking and to act as a window onto conceptual evolution.
3. A detailed elaboration of the nature of new knowledge and its relation during evolution to prior knowledge.
4. The relationship between new knowledge and the setting in which that knowledge is constructed.
5. A proposition that elaborates whether and how prior knowledge illuminates sense-making in unfamiliar settings.

Elements 1 and 3 consider the potential role for current mathematical knowledge in the development of new knowledge. Elements 2 and 4 consider the analogous role for external resources. Element 5 aims to present a coherent statement in which the respective role of both internal and external resources are discussed in relation to the situatedness of mathematical knowledge. These five elements will act as organizing points for the subsequent sections.

We will illustrate the five elements with reference to the developing ideas of a pair of children, part of a larger study of 10-11 year-old children's construction of resources for randomness (Pratt, 1998a, 1998b, 2000). At the risk of repeating ourselves, we should emphasize that the model we outline here is derived from a number of prior studies (some of which are collected in Noss & Hoyles, 1996) and, with respect to the issue of randomness discussed in this paper, from detailed data analysis of Pratt's (1998a) study. The data reported in this paper, therefore, should be viewed as exemplifying the theoretical construction, rather than as a corpus of data from which we are attempting to (over-)generalize.

The theoretical framework presented here has emerged through reflection on the data from the original study. Although in this paper we present fragments of that data chosen for its illustrative value, it is helpful for the reader to know some of the methodological details of the original study. (The computer environment itself will be described in Element 2.)

The methodology of the broader study was founded on an iterative design in which the study of the children alternated with the development of computer-based tools and resources. We worked with a new group of children in each iteration. 32 children, aged 10 and 11 years, were studied in all, with the final iteration involving 16 children. Each child was interviewed individually using a semi-structured schedule in which the children expressed their ideas about randomness before working with the computer-based tools. The early interactions with the tools provided further insights into the nature of this initial knowledge. Element 1 of our theoretical model represents insights gained from those interviews and early interactions with the microworld. The children worked with the microworld in pairs for between 2 and 2.5 hours. In these activities, the researcher acted as a participant observer, interacting with the children in order to probe the reasons behind their actions, later interpreting these reasons in the light of observations based on their and other children's work. The approach throughout was to allow the children to be in control of their explorations, making decisions and moving

in directions of their own choice. Most interventions were intended to be neutral with respect to changing the direction of the children's thinking. Typically such interventions were concerned with explaining technical matters or probing what lay behind a child's actions. More directed interventions were used when the children were clearly stuck or embarked upon a path with no potential pay-off from either research or learning perspectives. Experimental interventions were also used to probe for the maximal level of performance that the child's internal resources, supported by the computer's tools, could achieve. Such interventions were only used when a child seemed to be particularly confident and already performing with some degree of fluency.

The actions and discussions of the children were video-taped. Case accounts, avoiding as far as possible interpretation, were developed from the transcripts. (Some of the illustrative transcripts in this paper include descriptive comments from the case accounts.) These accounts were the basis for subsequent case analyses in which various inferences were made as to why and how the children's internal resources were modified. The case analyses drew on the transcripts of the pre-interviews as well as the case accounts of the activities. These case analyses were made available for a colleague and differences of opinion about the inferences were discussed and resolved.

A trace of each pair's work was developed. This trace indicated the path taken including the situated abstractions articulated during that journey. (The trace in Figure 5 is a graphical summary of what were very detailed traces.) The case analyses and traces were studied and compared across children to identify consistencies and marked differences. Consistency of an issue from case to case provided a check on reliability whilst variations suggested the limitations. The theoretical model developed here is based on this corpus of data although we only provide one case as an illustration of the theory.

Element 1: Naïve knowledge

In these data, children were observed to use four separable resources for articulating randomness, namely unsteerability, irregularity, unpredictability and fairness. In each case, we use a single example to illustrate how the resource was expressed more generally. In this way, we provide a first glimpse of our view of the nature of naïve mathematical knowledge.

(i) Unpredictability

Objects like dice were often described as unpredictable. 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.”

The apparent connection between unpredictability and random devices was very commonly expressed. Indeed, for many children, a way of checking the randomness of a device was to see if it was possible to predict outcomes.

(ii) Unsteerability

Phenomena were often seen as random when no known agent was involved in determining the result. 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.”

The primitiveness of unsteerability is also seen in the way it is often used to “explain” other resources for randomness. Very frequently, children associated unpredictability with unsteerability. On such occasions, unpredictability was usually seen as the outcome of uncontrolled input.

(iii) Irregularity

A third way of making sense of random phenomena was through reference to the lack of a regular pattern in sequences of results. This resource was often linked closely to prediction so that patterns were conjectured on the basis of past results and then used to make predictions, which were tested by further trials. 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

Fairness was often a defining characteristic of randomness. 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.”

The excerpts illustrate the nature of naïve knowledge of randomness as observed throughout our interviews and as frequently observable in early interactions with the software. Below we outline five aspects of this naïve knowledge.

We observed children articulating different naïve resources within moments of each other. 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.

It seemed that different aspects of the setting cued different resources. For example, the physical appearance of an everyday spinner sometimes cued an awareness of a lack of fairness, whereas in the computer setting, attention was drawn to the use of the strength control (see Figure 1) and so unsteerability and unpredictability were cued. The change in focus of attention for the child seemed to stimulate different resources.

Their interchangeability suggests that these naïve resources were connected in the sense that when one resource was triggered another (connected) resource was automatically available as well. It is important to have some sense of the strength of that connection within naïve resources. We often noticed children articulating different naïve resources as if they were interchangeable when to the expert they may have seemed contradictory. For example, a non-uniform spinner was often regarded as non-random because it was not fair, even though it was neither predictable nor controllable. Some children resolved this by expressing degrees of predictability (a uniform spinner being less predictable than a non-uniform one) but most children never articulated any awareness of such contradictions. If the fairness and unpredictability resources had been strongly connected, the triggering of fairness would have triggered unpredictability and one would have expected the children to recognize some dissonance. It seems then that the apparent interchangeability of naïve resources was not due to a strong connection between them, and may simply have been because attributes of the setting that triggered one resource often appeared coincidentally with aspects that triggered another.

Naïve resources were used, at least initially, to construct meanings for long-term as well as short-term behavior. There was little attempt in the early interactions with the gadgets to diagnose their behavior in terms of aggregated results. Even when relatively large numbers of trials were generated, the children searched for local regularities in results.

Our final observation regarding naïve knowledge of randomness relates to the relationship in children's minds between random and deterministic behavior. The short-term behavior of the gadgets was often explained by reference to cause and effect. A typical way of working was for the children to begin by conjecturing possible deterministic reasons for a gadget's initial behavior. Further results though did not support that conjecture and an alternative deterministic conjecture was found. When this failed too, yet another reason for the behavior was proposed, though at some point, when sufficient testing has been carried out, the search for such explanations came to an end and the behavior was described as random.

There is a strong sense in which stochastic meanings are constructed to deal with those aspects of the world that cannot be deterministically explained. It is no coincidence that the children's actions usually involved searching for deterministic behavior since it is impossible to look for stochastic behavior when it is characterized only in terms of the absence of various attributes: a lack of predictability, a lack of sequential patterns, and a lack of control (fairness is a notable exception being sometimes observable from the physical characteristics of the device). A richer appreciation of the stochastic must involve the identification of positive features and these lie in long-term behavior.

The above characterization of naïve resources for randomness is strikingly resonant with diSessa's, who describes p-prims as multitudinous, small pieces of knowledge that are self-evident, not needing justification and weakly connected. Naïve knowledge of randomness then seems to have many of the characteristics of diSessa's p-prims.

The children in this study hardly ever articulated in the initial interviews an appreciation of the significance of aggregating results over the long term. When asked how to determine the fairness of a dice, children referred to the appearance of the dice (fairness) or they threw the dice a relatively small number of times (irregularity). Children never referred

to the frequencies of a particular outcome over a period of time. If any aggregating strategy existed at all for these children, it had low cueing priority. We will see later how new knowledge of this kind did begin to emerge.

Element 2: The Setting

A detailed model of mathematical abstraction in context demands a clear description of the context itself, including the available tools and resources. We chose to use a computational setting because we believe that a computational environment provides a unique opportunity to tune the tools available within the setting to our research agenda. We iteratively designed a stochastic microworld, Chance-Maker, based upon Boxer,¹ which is particularly suited to our purposes. The aim was that this setting would enable us to observe children's naïve and emergent knowledge through their actions with the tools. In addition, we required that the virtual environment would enable us to look closely at the children's thinking as they moved from one setting to another. Let us consider exactly how Chance-Maker provided these three windows--on naïve knowledge, on emergent knowledge and on new knowledge across settings.

Providing a window on naïve knowledge

In order to observe naïve knowledge, we had to design tools that would cue current resources for randomness, unpredictability, unsteerability, irregularity and fairness. Chance-Maker contained a number of gadgets that behaved in many identifiable respects like their everyday counterparts. Consider for example the COIN gadget.

The COIN gadget can be activated using a "strength bar" causing it to spin much like its everyday equivalent. Figure 1 depicts the strength as a solid black bar with a circular switch at one end. Imagine the child controlling the strength by allowing a tube (the black bar) to fill with a red fluid until the switch is clicked. The strength of the throw, 30% in this case, is represented by the amount of red fluid. Clicking directly on the COIN causes it to be

tossed with the strength already indicated on the strength bar i.e. the strength used on the previous occasion. It spins like a coin, indicating the result on its front surface.

Figure 1

We hoped that asking children about the behavior of the COIN, and specifically whether they thought it was working properly, would provide a valuable window on how they thought about randomness.

Providing a window on new knowledge

In order to observe the evolution of new knowledge, we needed to provide tools that would perturb thinking. We did this by offering structuring resources that went beyond those to be found in everyday situations. For example, the COIN gadget can in fact be opened up to reveal a range of new tools (Figure 2).

Figure 2

Of particular importance is one tool that we encouraged the children to call the “workings box.” This tool captures the mathematical essence of how the COIN works. In one sense, the workings box is a formal symbolic description of the behavior of a coin. In this respect, it encapsulates the notion of a probability distribution. In addition, and we see this as a critical design attribute, the workings box can be edited by the child. When the workings box is modified, the behavior of the COIN will change correspondingly. We intended that the child’s re-programming of the workings box would provide a window on how they began to understand that correspondence.

We also provided a repeat tool (Logo-style), to facilitate the use of large numbers of trials. The ON/OFF button allows the graphics to be toggled essentially to save time when repeating many trials. The NEW button allows the start of a new experiment by clearing out old results; the results box simply lists all results since the NEW button was last clicked. By clicking on one corner of the results box, a chart of the results is displayed instead of the list.

Two types of chart are possible, generated by clicking on the PIE or PIC buttons. A pie chart shows the proportions of the results for each outcome whilst a pictogram indicates in absolute terms the frequency of each outcome. The pictogram is not automatically scaled and so differences between the number of results for each outcome are often brought to the child's attention. There is a feature to scale the pictogram; by increasing the scale, differences between rows in the pictogram will then be reduced proportionately. We intended that the use of the repeat and charting tools would provide a window on emerging ideas about aggregation of results and the relationship with the configuration of the workings box.

Asking the children to mend the COIN gadget provided a sense of purpose for their activity and, through their activity with the workings box, a window on their emerging knowledge about randomness.

Providing a window on different settings

We needed to create different settings to enable us to observe children's thinking during the transition from one setting to another. We replicated the features of the COIN gadget in a series of new gadgets. By keeping the structuring resources more or less invariant across the different gadgets, we could observe the influence of prior expectations, differences in appearance, and differences in detail, for example in the contents of the workings box. In this paper, we will focus on the children's engagement with just three gadgets, the COIN, the SPINNER and the DICE. Figure 3 illustrates the workings of the SPINNER gadget. The workings box and the graphic of the SPINNER indicates a bias towards 1. When the child modifies the workings box, the graphic automatically changes to remain consistent.

Figure 3

Figure 4 depicts the workings box of the DICE gadget, which contains a bias towards sixes.

Figure 4

By tracing children's actions and articulations across gadgets as they explored whether the gadget was working properly and as they attempted to mend a broken gadget, we were able to observe the relationship between the specificities of the setting and the construction of new knowledge. We were also able to observe whether and how knowledge constructed in a previous setting (a prior gadget) was utilized, or not, in making sense of activity in an unfamiliar setting (a new gadget).

Element 3: New knowledge

We are not proposing at this point to present the chronology of how new resources evolved, but, rather unconventionally, to set out in advance the nature of the new knowledge as it eventually emerged in order subsequently to focus on how that ending was reached. More specifically, in elements 4 and 5 we will trace the evolution of this knowledge in relation to the settings in which it was constructed, and at that point we will illustrate in much more detail the nature of those new resources. For now, we offer a definition of two new resources, but a proper appreciation of those definitions will emerge during elements 4 and 5. We have labeled the new resources Large Number (N) and Distribution (D).

The Large Number resource (N)

As children worked, they naturally used the graphing tools to display the frequencies (or in the case of the pie chart, the relative frequencies) of each possible outcome. They tried to use this graphical information to make sense of the gadget and in particular to decide whether further mending was necessary. Typically, they began using small numbers of trials to test the operation of the coin and spinner gadgets, but as time went by they worked with larger numbers of trials and began to observe invariants in the appearance of the pie chart. In situations where the workings box had been modified to represent a uniform distribution, they eventually recognized that the pie chart had roughly equal sectors when the number of trials was large. We regard this insight as intuitively related to the Law of Large Numbers,

and we label it the Large Number resource, N , which can be expressed as “the larger the number of trials, the more even the pie chart.”

The Distribution resource (D)

In order for N to develop in such a way that it can be applied successfully to non-uniform distributions, there is a need to recognize the formal description of distribution inherent in the workings box in the gadgets. This involved the construction of the second resource. Children played with many controls to identify the source of variation in results. Indeed the number of trials is one such control. In their search for powerful explanations and control over the behavior of the gadgets, children would typically discover or conjecture that the workings box could influence the appearance of the graphical output. Even when the number of trials was small, they would articulate expectations that changing the appearance of the workings box might have direct consequences for the appearance of the pie chart. We regard this insight as intuitively related to the concept of distribution. The Distribution resource, D , can be schematized as “the more frequent an outcome in the workings box, the larger its sector in the pie chart.”² Some children were able to relate N and D to each other, expressing the idea that “the more frequent an outcome in the workings box, the larger its sector in the pie chart, provided the number of trials is large”. We regard this type of articulation as indicative of coordination across N and D .

Let us summarize our argument to date. We have illustrated a model in which stochastic knowledge starts out as connected to a weakly coordinated set of primitive resources for randomness, abstracted from everyday experience, and with many of the attributes of p -prims. There appear to be few (if any) high-priority resources associated with long-term aggregated behavior. Tuning towards expertise can occur in the sense that two new resources, the Large Number and the Distribution resources might emerge through interaction with carefully designed external resources, and illustrative evidence for how these new

resources emerge will be presented in the next two sections. Further, we will illustrate the apparent coordination of N and D.

We now turn to elements 4 and 5, the relationship between the knowledge and the setting. The data in elements 4 and 5 is chronological and is offered as illustration for theory for the micro-evolution of mathematical knowledge. The story begins with two children, Anne and Rebecca, both age 10, engaging with the gadgets at top level i.e. not using the mending tools that reside inside each gadget. The story develops as they first play with the COIN gadget, then the SPINNER and finally the DICE gadget.

Element 4: The role of setting

Anne and Rebecca at top level

We have already illustrated (in element 1) the nature of naïve knowledge and (in element 3) the Large Number, N, and the Distribution, D, resources. We now wish to trace the cueing and construction of N and D alongside that of more naïve knowledge and the use of external resources.

Figure 5 summarizes the data in the form of a trace, which charts the evolution of the internal resources as articulated by Anne and Rebecca. Figure 5 displays the approximate times when Anne and Rebecca articulated particular internal resources. The “events” indicated on the trace often spanned several minutes; the positioning of the resources on the trace gives a rough indication of the timing. On some occasions these resources were expressed as conjectures, “perhaps the strength controls the number on the dice,” whereas on others the children seemed to be proposing a conclusion, “the more often you throw the dice, the more even the pie chart gets.” It is not always easy to make this distinction. For our purposes, the distinction is not critical, since, whether conjecturing or forming a conclusion, the same internal resource is being cued.

Figure 5

In Figure 5, we see many events in the first section of the trace where naïve knowledge was cued. During this first period, Anne and Rebecca were asked to play with the gadgets at top level, partly to gain a feel for the gadgets but also to form some initial views as to which gadgets seemed to be working properly and which might be broken. This play was at top level; they had not yet been introduced to the tools within each gadget. The main control available to Anne and Rebecca was the strength control. In playing with the coin, Anne said that she was testing it³:

1. Anne: . . . to see if it comes up tails again.

We can interpret this comment either as testing the COIN's irregularity, based on a perceived pattern in previous results, or as testing its unpredictability as shown on the trace. A few moments later, the girls were proposing that the computer was exerting control, a special case of steerability.

2. Dave: Do you think the COIN is random?
3. Rebecca: Not really . . . it's probably been programmed to do it, in a loop.
4. Anne: I don't know.
5. Dave: What do you mean by "in a loop?"
6. Rebecca: Well, it's programmed to do heads, then maybe heads again and then tails.
7. Dave: In some sort of pattern?
8. Rebecca: Yes.
9. Dave: Did you pick up some sort of pattern?
10. Rebecca: Not really.

In line 2, Rebecca is dealing directly with the task of deciding which gadgets are working properly. Rebecca's argument is that the COIN is not random because it is programmed. (Rebecca had a fairly sophisticated appreciation of the operation of computers.) We take this

as further support of how Rebecca regarded unsteerability as a facet of how a coin should behave.

When they moved onto the SPINNER, Anne's fairness resource was cued (line 11).

11. Anne: I think it might be a bit unfair because the one is much bigger than the others . . . The two, three, four and five is much smaller than the one, so one has more of a chance.

The unsteerability resource was also apparently cued since the girls tried to test whether the size of the strength control influenced the outcome (line 12).

12. Rebecca: I don't think you can really control it, because if you do it slowly, you never know, it could come on a three, if it starts off on a different number.

Rebecca felt at that time that the strength did not seem to control the outcome and instead conjectured that the outcome might be connected with the position of the number on the spinner. Anne and Rebecca frequently returned throughout the interview to similarly deterministic explanations of phenomena, which they could not find a way of describing otherwise.

We believe that the prevalence of conjectures and conclusions based on naïve knowledge during this preliminary stage can be accounted for in two main ways:

- (i) Anne and Rebecca described phenomena in terms of their naïve resources, since these were of much higher cueing priority than resources associated with longer-term aggregated behavior (or more extremely the girls may, as yet, have had no access to such resources for aggregated behavior).
- (ii) The strength control tool available at top level cued the use of unsteerability and unpredictability. In effect, the natural thing to do was to play with the strength control and test out ideas based on it. In contrast, there was no specific encouragement at top level to record or aggregate results.

Anne and Rebecca inside the COIN

When Anne and Rebecca moved onto work with the tools within the COIN gadget, similar p-prims were cued. In addition, the availability of the results box cued irregularity by offering a means of recording results. This period of exploration allowed Anne and Rebecca to try out conjectures related to the effect of strength. They found that none of these resources provided a consistent explanation of the COIN's behavior, as illustrated by this extract in which the girls had been conjecturing that there was a regularity to the COIN's results.

13. Dave: What do you feel, Rebecca, about what you were saying before that it would work to a pattern? Do you think there is a pattern?

14. Anne: Oh yes, I can see a bit of a pattern. Because that's got head, head, head, tail, and it's got . . . oh, where was it? I can't think . . . oh yes, it's got a head, oh no, that's not right, it's got tail, tail, tail, tail, that time.

15. Rebecca: I'm not too sure.

16. Anne: I don't think there's much of a pattern really.

17. Dave: Did you expect there to be a pattern from what you were saying before?

18. Rebecca: Yes, I did.

The availability of the pie chart led to Anne and Rebecca to consider aggregated patterns. The speed with which the computer generated results also made it easy for them to consider large-scale experiments. Even so, there was some reluctance to do so. Because Anne and Rebecca did not start new experiments, there was a gradual accumulation of results, and a tendency to increase the number of trials each time as confidence grew. At each stage, the charts were consulted and comments were made about the evenness or otherwise of the pie chart or pictogram. This process resulted in a recognition that the pie chart became "more even" as more trials were included in the experiment, the first articulation of N , after twenty minutes of trying to make sense of the COIN gadget.

19. Anne: I think it's the highest the number, the even more it gets.
20. Dave: The higher the number, the more even it gets? (Dave reverses the two words "more" and "even.")
21. Anne: Yes.
22. Dave: Do you agree, Rebecca?
23. Rebecca: Because the other time, when we did less numbers, it was half um . . . even really.
24. Dave: Do you agree with that--the more times you do it, the more even it's getting?
25. Rebecca: Yes, it seems to be.

The context made it clear that by "number" in line 19, Anne meant the number of trials in the experiment.

We believe that the Anne and Rebecca articulated N at this time because:

- (i) The pie chart and the pictogram focused attention on aggregated results.
- (ii) The design of the Chance-Maker environment made accumulation of results natural and the researcher's interventions gave such activity positive feedback.
- (iii) Resources like unpredictability were insufficient to explain the emerging predictability of the pie chart.

Anne and Rebecca inside the SPINNER

When Anne and Rebecca began to work with the SPINNER gadget, one might have expected them to begin by using N. In fact, this did not happen. The "unfair" appearance of the SPINNER caught their attention. The numbered lines up to this point have all referred directly to transcripts of discussions with or between Anne and Rebecca. In places, such as line 29, we have also numbered descriptive notations of the videotape, where these are direct quotations from our own translations of the tape.

26. Anne: Oh, look, it's got to choose from different numbers.

27. Rebecca: That's definitely more chance of it landing on one then.
28. Anne: Yes . . . There's more chance of it because that's (pointing to the 1 sector) much bigger than these (pointing to the other four sectors).
29. They make the strength 50 and do 50 trials. The pictogram shows most ones, then fours with threes the least. The pie chart confirms this.
30. Anne: It's a bit like that really (pointing to the spinner). I thought it would be more because one's more of a chance of getting there because it is much bigger.
31. Dave: Is the SPINNER random?
32. Rebecca: Not really . . . because there is more chance of it landing on one, than the other numbers.

Anne and Rebecca wanted to make the SPINNER fair and the only way to do this was through the workings box. For the first time, the workings box became a focus of their activity. They edited the workings to read: `choose-from [1 2 3 4 5]`. However, they only chose to carry out 50 trials and the pie chart did not appear fair, showing a larger 5 sector.

Dave asked the girls how the pie chart might be made to look fair expecting to cue N.

33. Dave: If your aim was to make that pie chart look more even, what would you do?
34. Anne: I'd make the five a bit smaller. . .
35. Rebecca: I'd make the others a bit bigger.
36. They begin to edit the workings.
37. Anne: Why don't you put them all the same number? That would be even then.
Like put three on one, three on two. That would be fair because that would be even then.

The question (line 33) was answered first in terms of the pie chart itself and then in terms of the workings box. Anne's response showed that she understood that the workings box could influence the appearance of the pie chart (and thus is an articulation of D) but she

did not have a proportional view. In fact, the girls settled on a workings box that read: choose-from [1 1 2 2 3 3 4 4 5 5]. They again used 50 trials as a new experiment suggesting that N was still not being cued. The pie chart was not fair, displaying most ones, provoking this reaction:

38. Rebecca: I wonder what would happen if we took one more away?

Again Rebecca showed that she appreciated a link between the workings box and the pie chart but she did not yet recognize that they were using too few trials. They went on to edit the workings to read: choose-from [1 2 2 3 3 4 4 5 5]. After 50 trials, the pie chart showed most threes. They edited the workings back to read: choose-from [1 2 3 4 5] but 50 more trials still showed a bias, now towards 1s.

39. Anne: It's a tiny bit even but there are more ones.

40. Rebecca: Yes, definitely more ones.

41. Anne: Yes, but it's fairly even, if you see what I mean.

42. Dave: So, if you are trying to make this pie chart more even what could you do?

43. Rebecca: Maybe throw it more times like we did with the coin?

44. Anne: Yes.

For the first time during the work with the SPINNER, N had been cued. Over the next few minutes, Anne and Rebecca tried increasing numbers of trials and found that the pie chart came out more and more fair.

45. Rebecca: There's a higher number, so the more chance of it being even, I think . . .

The more times you throw it, the evener it seems to get. And I think that's because there's more chances for a number to come up than if you do it say fifty times.

Anne and Rebecca had been working with the SPINNER intensively for nearly 15 minutes before N was cued even though they had articulated N at the end of working with the COIN gadget. We make the following observations about Anne and Rebecca's activity:

- N, as constructed during the COIN activity, was available as a resource for sense-making as evidenced towards the end of the SPINNER activity.
- N, during the early part of the SPINNER activity, had low cueing priority as shown by the considerable delay before it was evoked as a resource.
- N had low cueing priority at this time because the SPINNER's appearance cued the fairness p-prim, in effect blocking the cueing of N.
- The need to make the SPINNER fair encouraged the use of the workings box, which led to the construction of D. The focus on seeing the workings box as a control distracted the girls from considering the number of trials as a control.
- N was eventually cued because neither D nor any part of their naïve knowledge could satisfactorily explain the SPINNER's unfair pie chart.

The unfolding of the above activity structure reveals how the children brought their naïve internal resources to bear towards their goal of making the spinner work properly. The setting allowed the children to explore the gadgets with some confidence to search for patterns of behavior. The interaction of internal and external resources allowed the children to construct new internal resources to make sense of the COIN's behavior.

Element 5: Unfamiliar settings

Anne and Rebecca inside the DICE

The trace in Figure 5 shows a marked difference in the nature of the internal resources cued during DICE activity and during SPINNER activity. There is relatively little evidence of the use of naïve knowledge (though it has not, of course, completely "disappeared"). Now, N is cued quickly and D is cued soon after. In fact, Anne and Rebecca began using 1000 trials from the start of their work with the DICE gadget.

46. Dave: What's the advantage of doing it a thousand times?

47. Anne: You get more and you can sort of estimate.

48. Dave: What do you think will happen to the pie chart?

49. Anne: Maybe a bit even.

50. Rebecca: Maybe it's going to be even again because it seems to go more even the more times you throw it.

51. Anne: I think it's the more you throw it, the more even it gets.

52. Rebecca: Yes, that seems to be the case.

53. Anne: Because that's what happened most of the times, the more you get, the more even you get.

Here, N was used inappropriately; they assumed the unfairness in the pie chart could be smoothed out by applying more trials. Having proved to have explanatory power in both the COIN and the SPINNER, N has taken on high priority, thus dominating the recently constructed D. Anne and Rebecca also have available the fairness resource, which perhaps makes them predisposed to an expectation that the pie chart should be fair. The mutually supportive nature of N and the fairness resource possibly raise the cueing priority of both.

When the pie chart showed too many sixes, Anne and Rebecca's attention was drawn to the workings box.

54. Anne: I think sixes is popular because there's quite a lot of sixes in the choose-from (pointing to the workings box) . . . There might be a lot more sixes so it's got more chances of getting more sixes on it.

A few minutes after beginning to work with the DICE gadget, Anne articulated the influence of the workings box in the sense that the bias towards the sixes was causing the pie chart to look non-uniform. As a result, Anne and Rebecca began to edit the workings box and, after considering choose-from [1 1 2 2 3 3 4 4 5 5 6 6], they settled on choose-from [1 2 3 4 5 6]. They carried out 1000 trials and the pie chart appeared fair.

Dave wished to probe how Anne and Rebecca understood this result and the earlier one when the DICE had been biased towards sixes.

55. Dave: So, doing it lots of times, like a thousand, has made it more even this time, but, before we altered it when the workings were as before, it didn't make it more even. So, what do you think it did do?

56. Rebecca: Because there's more chance of getting a six. When it stops it might land on a six.

57. Anne: Because the workings were unfair.

Now, about 13 minutes after beginning to work with the DICE gadget, Dave probed Anne and Rebecca's emerging sense of the role of the workings box.

58. Dave: Let's say we were playing a game, and for some peculiar reason in this game, it would have to be a computer game because we are using the computer dice, we wanted there to be a good chance of getting ones, and a fairly good chance of getting twos but a pretty low chance of getting anything else. It's a strange game. How would we make this dice behave like that?

59. Rebecca: You have to put more of the numbers on here.

60. She begins to edit the workings until they read:

`choose-from [1 1 2 2 3 4 5 6]`. They test this out by repeating 1000 trials.

61. Dave: What will the pie chart look like?

62. Rebecca: More twos, more ones and less of the others.

63. Dave: How will the ones and twos compare?

64. Anne: Roughly even.

Anne and Rebecca were now recognizing limitations of N, which had not been apparent when they began working with the DICE gadget. In other words, they seemed to have begun to construct a relationship between N and D.

We now wish to illustrate how N and D can become coordinated. Anne and Rebecca were able to discuss how the pie chart's appearance reflected that of the workings box, provided the number of trials was high; we will refer to the coordinated resource as N-D. A corollary to N is that small numbers of trials are likely to generate non-uniform pie charts. When the two resources are coordinated properly, the corollary can be articulated as a suggestion that the pie chart may not reflect the appearance of the workings box when the number of trials is low. Anne and Rebecca expressed this idea towards the end of their period of working with the DICE.

65. Dave: If we had done it only fifty times, instead of a thousand, do you think we would have got a picture like that?

66. Anne and Rebecca (simultaneously): No.

67. Anne: A bit more uneven.

68. Rebecca: There'd be more ones and twos.

69. Anne: They'd probably be about even.

70. Rebecca: Maybe.

71. Dave: 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.

72. Rebecca immediately edits the workings box to read:

```
choose-from [1 1 1 2 2 3 4 5 6].
```

Anne's comment (line 67) seems to indicate a coordination of N and D, though subsequently (lines 68-70) Anne and Rebecca may still have been cueing N inappropriately.

We make the following observations about Anne and Rebecca's activity at this point:

- N had taken on higher cueing priority as a result of it having been found to be a reliable sense-making resource for two different contexts, the COIN and the SPINNER
- N was also likely to be cued because of its consistency with the fairness resource (even though it was inappropriate to apply fairness to long-term behavior in this way).

- Less initial stress was placed on the workings box than had been the case with the SPINNER because there was no visual cue of unfairness in the appearance of the DICE.

We believe that the coordination of N and D took place at this time because neither individually was sufficient to “explain” the unfairness of the pie chart for an unfair workings box.

Discussion

We summarize the findings as follows. The children (exemplified by Anne and Rebecca) came to the computer-based interviews with access to four naïve resources, which we have called unpredictability, unsteerability, irregularity and fairness. These resources, similar to those employed by diSessa to describe p-prims in his model of conceptual change, dominated what children said throughout the interviews, although gradually--as new resources for long-term aggregated behavior emerged--they became restricted in use to describing short term behavior.

During the interviews, we noted the emergence of two new resources: Large Number, N, and Distribution, D, which differed from the initial resources, in that:

- they were associated with a mathematically coherent structure for randomness;
- they were not isolated nor weakly connected but offered mutual support through their coordination;
- they included within their articulation a consequence e.g. more trials leads to the consequence of a more even pie chart.

N and D represented abstractions of long-term behavior, a sense of an invariant relation that connected the number of trials, the configuration of the workings box and the appearance of the pie chart. At the same, the articulation of N and D was bounded strongly within the linguistic and material resources available within the activity. N and D are situated abstractions, initially having low cueing priority, evidenced from the continued reliance on

naïve resources by Anne and Rebecca in moving from the COIN to the SPINNER gadgets. But as the task-interviews progressed, N and D took on higher cueing priority, until these became mutually coordinated. By coordinated here, we mean that by the end of the work with the DICE gadget, both resources were being used together as mutual support rather than in isolation--the graphical representation of the results was now seen as contingent on both the number of trials and the configuration of the distribution in the workings box.

A model for the construction of situated abstraction

We claimed earlier that a model of the micro-evolution of mathematical knowledge required five elements. We claim some progress in respect of elaborating these elements at the required level of detail which will offer further illumination on our three fundamental questions: (i) What is abstracted? (ii) How do abstractions evolve? (iii) Is mathematical abstraction de-contextualized? We believe that to attempt to answer these questions we must connect macro and micro models: our approach is to relate our macro construct of situated abstraction to the micro construct of p-prim.

For the first element, we sought a clear description of the nature of naïve mathematical knowledge including the grain size of emergent mathematical knowledge. The naïve resources for randomness, unpredictability, unsteerability, irregularity and fairness, operate as a means of contextualizing phenomena. That is to say, the child identifies an aspect of a phenomenon, which triggers the unpredictability resource, say, enabling the triggering of other connected naïve resources such as unsteerability. Piaget and Inhelder (1951) have argued that the concept of randomness first emerges because, in contrast to all previous experiences, such phenomena are not explainable in terms of reversible operations. According to this view, random phenomena are fundamentally different from previously encountered situations and remain as such until the individual accommodates the concept of

probability. We suppose that contextualizing allows the organism to partition phenomena that do not conform to an operational structure.

These naïve resources for randomness do not have consequences and as a result have limited value for making sense of phenomena. In contrast, the typical p-prim enfold a class of implications (for example, Ohm's p-prim: more effort implies more results). In common with p-prim, we believe that naïve resources for randomness are abstracted directly from experience (playing with coins, dice and so on). Each becomes linked to aspects of the game-playing situation in such a way that specificities (quite possibly superficial) of the environment trigger particular resources. Hence, the appearance of a dice might trigger fairness at one point in a game whereas a focus on how the player is throwing the spinner or dice might trigger unsteerability. The apparent structuring of these resources (in the sense that when one is cued so might another) is based mainly on their coincidental appearance in the same situation (such as game playing). Any such structuring is weak because of a lack of justification, which might suggest that two coincidental resources are in fact contradictory. The nature of what is abstracted changes through tuning towards mathematical expertise as we shall see in revisiting elements 3 to 5 of the micro-evolutionary model.

Before then, we must consider the second element of the model. We described a setting that encouraged children to make conjectures and provided tools with which to express ideas and test out conjectures. The setting was designed to offer a range of situations in which mathematical resources for the aggregated behavior of random gadgets might prove to have explicative power. The Boxer environment was designed to provide a window for the researchers to observe the micro-evolution of knowledge about randomness, but at the same time to enable the children to scrutinize their own knowledge. We do not wish to suggest that the Chance-Maker environment itself (or any computer-based environment for that matter) has independent power to change children's thinking. On the contrary, other factors in the

setting contributed in significant ways. The task of mending the gadgets was important as it pointed the children's activity towards sense-making, and provided a point of reference for them in deciding whether that activity was complete. The questions and support from the researchers were sometimes critical in encouraging reflection and supporting them through technical difficulties. Nevertheless, Chance-Maker provided a framework in which this type of activity could naturally take place across contexts that differed in controlled ways, but contained certain common structures, such as the workings box and the graphing tools. We claim that the detailed elaboration of the setting, including the role of Chance-Maker, provides the second element of the model.

For the third element of the model, we sought a detailed elaboration of the nature of new knowledge and its relation during evolution to prior knowledge. At the macro-level, we are able to refer to the children's articulations, and it is these linguistic and physical actions that betray a dependence upon the structuring resources of the situation, and indicate the emergence of a situated abstraction. Linguistically, the children typically expressed their ideas in terms of the tools and resources available within Chance-Maker. More revealing still is the way that Anne and Rebecca appeared to reinvent, rather than simply re-use, the previously constructed situated abstraction, N, when moving from the COIN to the SPINNER. For them at least, N appeared to have narrow scope, at least in the first instance. Schematically we think of a situated abstraction as surrounded by a contextual neighborhood that describes the essential conditions, purposes, and features under which the situated abstraction was constructed. Recognition of the characteristics of the contextual neighborhood by the individual assists the identification of similar conditions under which the situated abstraction is triggered. In this sense, the idea of contextual neighborhood seems to be closely associated with diSessa's micro-level notion of cueing priority.

We would like to share a useful metaphor we use for thinking about the relationship between the macro and micro-levels. We observe children through their actions and their words at the macro-level. We hear their explanations of phenomenological behavior. On analysis, we classify some of these articulations as situated abstractions. But now imagine applying a microscope to these situated abstractions. What we would see, according to this metaphor, is a collection of linked p-prims in the process of being triggered. We believe there are examples where we can identify the p-prims that underlie Anne and Rebecca's articulations.

On many occasions, Anne (lines 19, 47, 51, 53) and Rebecca (lines 23, 43, 50) articulate the N situated abstraction, underlying which we recognize an idea of "more begets more", a precisely equivalent structure to the Ohm's p-prim. Thus there seems, in this case at least, a simple mapping between situated abstractions at the macro-level and p-prims at the micro-level.

More generally, it seems that the experience of working within the Chance-Maker setting allowed children to construct the N and D situated abstractions and that these articulations at the macro-level may signal new connections between the naive resources of randomness with existing or new p-prims. For example, a link between the unpredictability resource and the p-prims underlying the N situated abstraction would imply that, whilst a phenomenon may be recognized as random, it might still possess some degree of predictability in the long term. Furthermore, the co-ordination of N and D may have signaled further processing at the micro-level, the connecting of p-prims "beneath" the two situated abstractions.

For the fourth element of the model, we demanded an understanding of the relationship between new knowledge and the setting in which that knowledge is constructed. At the macro-level, Chance-Maker encouraged the construction of situated abstractions. In the study

reported above, randomness was instantiated as a manipulable computational system, so that children could see and act on the mechanisms of objects such as dice and coins. In other words, we built a system that phenomenonized the abstract, making it concrete. This is the key affordance of mathematical microworlds. A continuity of structuring resources from situation to situation (for example, the role and appearance of the workings box was common across the gadgets), though not necessarily immediately recognized by children, offered the opportunity for situated abstractions involving those structuring resources to have explanatory power beyond specific situations. At the same time, mathematically redundant tools (such as the strength control) built into the design allowed children to recognize that redundancy. In other words, the setting was designed to optimize tuning towards expertise by giving maximal support to expert-like resources.

Our analysis of children's articulations across situations enables us to elaborate the fifth element of the model, illuminating the degree of dependence on context. As part of tuning towards expertise, knowledge within a contextual neighborhood becomes less tied into the specific context of its genesis. Given new situations which have some similar features to those already experienced, and particularly in an environment--like Chance-Maker--in which prior resources can be road-tested, the contextual neighborhood can evolve and broaden into a domain of validity that describes the variety of circumstances in which a p-prim might be triggered. In other words, knowledge cannot escape from context, but the range of contexts in which it might be cued can and does broaden, as new connections are made. We saw how, in the study, the children began with a narrowly defined contextual neighborhood for the Large Number situated abstraction, N, applied precisely to the COIN gadget. When they began work with the SPINNER, there was no reason to suppose that N applied, or to put it more succinctly, that the phenomena observed while playing with SPINNER lay within the contextual neighborhood defined by COIN. Anne and Rebecca found that the resources they

had to hand were insufficient for making sense of the COIN, and used N as a way of overcoming the difficulty. In doing so, they found that it worked for both the COIN and the SPINNER: the contextual neighborhood for N expanded, with the result that N was called upon relatively quickly when Anne and Rebecca began to work with the DICE gadget.

Concluding remarks

We have described in some detail five elements that constitute a model for the micro-evolution of mathematical knowledge in the case of randomness. There are obvious limitations to the generalizability of our data, not least that the analysis emerges from reflection on data from just 32 children, and is illustrated here through the activity of just two. Despite these limitations, we believe that our model may have some explanatory and predictive power for making sense of the evolution of children's ideas of randomness and probability and may perhaps offer a framework for understanding the construction of mathematical meanings more generally.

From the point of view of randomness and probability, one of the striking findings of the literature is the inconsistent nature of student responses to stochastic situations (for example, Konold et al., 1993), responses that appear to be sensitive to the situation (Nisbett et al., 1983). Searching for a model of learning about the stochastic has, as a result, been problematic. The proposed model in this paper provides a micro-level explanation for such inconsistency, in the sense that naïve resources for randomness only serve to contextualize randomness and have no power to offer rational explanations of long-term behavior. At the same time, the model suggests that settings with particular properties (e.g., those attributes of Chance-Maker highlighted above) would support the construction of more expert-like, but situationally dependent abstractions about aggregated behavior.

It is reasonable to ask whether the types of insights gained by children from working with Chance-Maker are transferable to domains that are not computationally mediated. The

abstractions, N and D, are situated in the sense that they are constituted within the Chance-Maker context and derive their immediate meaning from that context. In this sense, the notion of situated abstraction is entirely consistent with Lave's hypothesis that knowledge is situated and the notion of transfer is an irrelevance. However, through reflection on children's attempts to make sense of new gadgets within the microworld, we predict in our model what will happen when children explore situations beyond the confines of Chance-Maker. We regard those situations as new domains for sense-making, just as the DICE was a new domain beyond the spinner, and the SPINNER was a new domain after the COIN. N and D live in the minds of those children. Some of the 32 original children will have organized N and D with greater priorities than others. They may well call upon N and D more readily than other children in situations where they are trying to make sense of stochastic phenomena. If the new situation is superficially different from that experienced in Chance-Maker then naïve resources, such as unpredictability, may be triggered. These resources will only be rejected if they fail to have explanatory power in those new situations. Real world situations often fail to provide feedback that indicates such failure, and under those circumstances N and D may remain dormant.

Of course, treating the world as just another domain of abstraction is highly speculative, and is not – at least on the basis of the study reported here – justified by the corpus of data on which the model is grounded⁴. Nevertheless, the micro-level mechanism at work in the model for generating and coordinating situated abstractions do not, on the face of it, appear necessarily to be tied to the specificities of the domain we constructed, except, of course, that we designed it as carefully as we could to maximize the probability of conceptual change.

Our final word briefly outlines a possible framework for applying the model beyond the case of randomness, and to revisit the question of context with which we began. We propose

that de-contextualization and contextualization might usefully be viewed not as antithetical but as different perspectives on the same micro-evolution of mathematical knowledge. Meanings--internal resources--are created by experiences, which can evolve from webbing connections between mental states as much as between a mental state and the physical world. These resources remain rooted in actual experience. Both naïve resources and situated abstractions are abstracted through making connections at the micro-level, the latter with consequences. So far so good. But what of mathematics as viewed by mathematicians, in which mathematical meaning appears to depend on being cut loose from context? We believe that what appears to the mathematician as an abstraction away from, and independent of, context, is in fact a broadening of contextual neighborhood. It is this model of abstraction taking place within, rather than outside, a context that elaborates the notion of situated abstraction. De-contextualization is a post-hoc perspective on mathematical knowledge that does not necessarily assist in understanding the trajectories of learning, which, according to our model, emerges out of a broadening of, rather than a cutting away from, context.

On the other hand, it is unhelpful to offer a perspective in which mathematical knowledge is trapped within a situation. When knowledge is constructed, new ideas may gradually gain higher priority within that domain but will continue to possess low priority in unfamiliar settings. When environments are made available in which those same ideas prove to be the most powerful sense-makers of activity within that environment, our model would suggest that these apparently situated meanings will take on higher priority within the unfamiliar setting.

Acknowledgement

We would like to thank Andrea diSessa for his careful and insightful comments on an earlier draft of this paper.

References

- Confrey, J. & Costa, S. (1996). A Critique of the Selection of "Mathematical Objects" as a Central Metaphor for Advanced Mathematical Thinking. International Journal of Computers for Mathematical Learning, **1**(2), 139-168.
- diSessa, A. (1983). Phenomenology and the evolution of intuition. In D. Genter & A. Stevens (Eds.), Mental models. Hillsdale, NJ: Erlbaum.
- diSessa, A. (1993). Towards an epistemology of physics. Cognition and Instruction, **10**, 2/3, 105-226.
- Dubinsky, E. (1991). Reflective Abstraction in Advanced Mathematical Thinking. In D. Tall (Ed.). Advanced Mathematical Thinking. Dordrecht: Kluwer Academic Publishers.
- Gray, E. M., & Tall, D. O. (1994). Duality, ambiguity and flexibility: a proceptual view of simple arithmetic. Journal for Research in Mathematics Education, **26**, 115-141.
- Hoyles, C. & Noss, R. (1992). Looking back and looking forward. In: Hoyles C. & Noss R. (Eds.) Learning Mathematics and Logo. pp. 431-468. Cambridge: MIT Press.
- Hoyles, C., Noss, R., and Pozzi, S. (2001) Proportional reasoning in nursing practice. Journal for Research in Mathematics Education. **32**, 1, 4-27.
- Konold, C., Pollatsek, A., Well, A., Lohmeier, J., & Lipson, A. (1993). Inconsistencies in students' reasoning about probability. Journal for Research in Mathematics Education, **24**, 392-414.
- Lave, J. (1988). Cognition in practice. Cambridge, UK: Cambridge University Press.
- Nemirovsky, R. (in press). On guessing the essential thing. In K. Gravemeijer, R. Lehrer, B. Van Oers, & L. Verschaffel (Eds.). Symbolizing, modeling and tool use in mathematics education. Dordrecht, The Netherlands: Kluwer.
- Nisbett, R., Krantz, D., Jepson, C., & Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. Psychological Review, **90**(4), 339-363.

Noss, R., & Hoyles, C. (1996). Windows on mathematical meanings: learning cultures and computers. Dordrecht: Kluwer.

Noss, R., Hoyles, C., & Pozzi, S, (in press). Abstraction in expertise: a study of nurses' conceptions of concentration. Journal for Research in Mathematics Education.

Papert, S. (1982). Mindstorms: Children, computers and powerful ideas. London: Harvester Press.

Piaget, J., & Inhelder, B. (1951). The origin of the idea of chance in children. [by] Jean Piaget and Bärbel Inhelder; translated [from the French] by Lowell Leake Jr, Paul Burrell and Harold D. Fishbein. New York: Norton.

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

Resnick, L. (1991). Shared cognition: thinking as a social practice. In L. Resnick, J. Levine, & S. Teasley (Eds.), Perspectives on socially shared cognition, (pp. 1-20). American Psychological Association.

Sfard, A. (1991). On the dual nature of mathematical conceptions: Reflections on process and objects as different sides of the same coin. Educational Studies in Mathematics, 22, 1-36.

Sfard, A. (1994). Reification as the birth of metaphor. For the Learning of Mathematics, 14(1), 44-55.

Wilensky, U. (1991). Abstract meditations on the concrete and concrete implications for mathematics education. In I. Harel & S. Papert (Eds.), Constructionism (pp. 193-204). Norwood, New Jersey: Ablex.

¹ Andy diSessa heads a team at Berkeley, which is developing a “computational medium,” named Boxer. This project is seen as extending the notion of literacy to a new domain, where users express themselves in various ways, including mathematically, in various modalities (graphic, literal, computational . . .).

² This schematisation places no emphasis on proportion, which is probably a fair reflection of the how the D resource was articulated. A few children discovered that a workings box that read: `choose-from [1 2 3 4 5 6]` was essentially the same as one which read: `choose-from [1 1 2 2 3 3 4 4 5 5 6 6]`. Most children however articulated D in terms of the frequency of an outcome rather than its proportion.

³ We have re-numbered the lines from our case accounts in a way that is convenient for this paper. The excerpts are taken from various sections of the data and, even though different excerpts are numbered consecutively in this paper, they did not necessarily occur immediately one after the other in real time.

⁴ We have, however, made some progress in this direction in the context of studying mathematical knowledge in workplace contexts: see, for example, Hoyles, Noss and Pozzi (2001); Noss, Hoyles and Pozzi (in press).

Figure captions




Figure 1. The COIN gadget

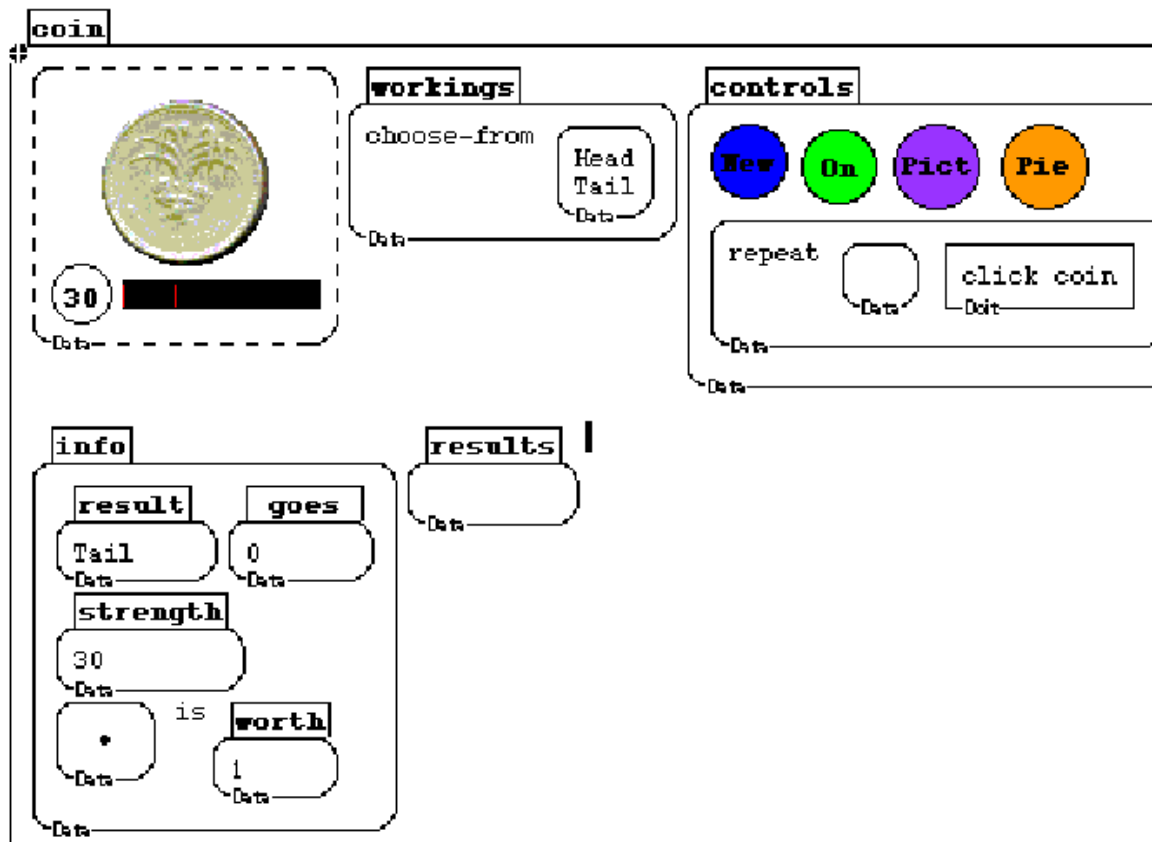


Figure 2. The tools in the COIN gadget

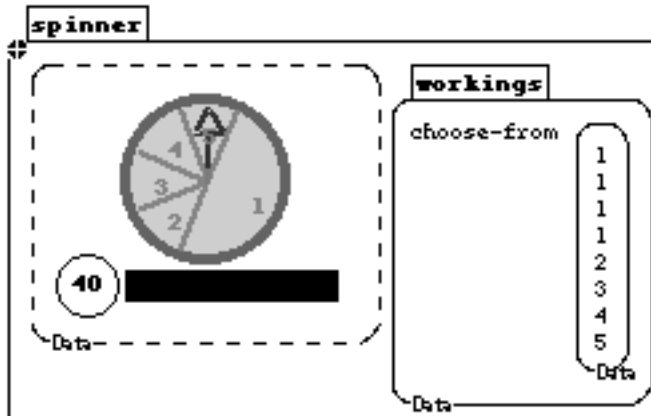


Figure 3. The workings inside the SPINNER gadget

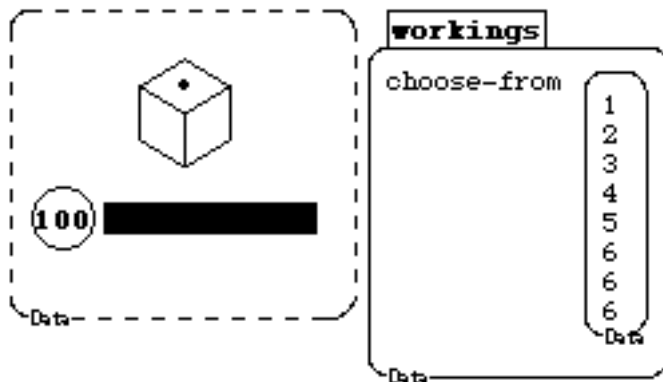


Figure 4. The workings inside the DICE gadget




Figure 5. A time trace showing when Anne and Rebecca articulated various resources

Light circle = primitive resource; Bold circle = new resource; Filled circle = coordinated resource.

P = unpredictability; S = unsteerability; F = unfairness.

I = irregularity; N = large number; D = distribution.