

Abrahamson, D. (2006, April). Bottom-up stats: Toward an agent-based “unified” probability and statistics. In D. Abrahamson (Org.), U. Wilensky (Chair), and M. Eisenberg (Discussant), *Small steps for agents... giant steps for students?: Learning with agent-based models*. Symposium conducted at the annual meeting of the American Educational Research Association, San Francisco, CA.

Bottom-Up Stats: Toward an Agent-Based “Unified” Probability and Statistics
 Dor Abrahamson
 Graduate School of Education
 University of California, Berkeley

Abstract

Computational environments can do more than display old ideas in a new medium. Building on Wilensky and Papert (Wilensky, 2006; Wilensky & Papert, 2006), I examine multi-agent modeling-and-simulation technology that both restructures old ideas in probability and statistics and illuminates connections between these domains of practice. I first present two computer-based activities, one in probability and one in statistics (ProbLab, Abrahamson & Wilensky, 2002), and then compare between them in terms of perceptual, procedural, and conceptual dimensions. I demonstrate how ‘student-with-computer’ overlapping experiences in the probability and the statistics simulations—in terms of obtaining samples, in terms of the temporal dimension of these acts, and in terms of visual metaphors apt for thinking about either activity—suggest a restructuring of probability-and-statistics practices as nuanced epistemological variants governed by the parameter ‘experimentation vs. exploration.’ I implement this *restructuring* (Wilensky & Papert, 2006) as a proposed design object—the *Platonic combinatorial space*, a “population” from which random compound-event outcomes are “drawn” in probability experiments. Finally, I present a “bottom-up statistics” explorative conceptual model that builds a Law-of-Large-Numbers explanation for the Central Limit Theorem.

Introduction

Consider a box full of marbles: green marbles and blue marbles. You dip into this box a device that scoops out samples of exactly four marbles per scoop. You make a note of how many of the four sampled marbles came out green, dump the marbles back in, and repeat this activity several times. Now, further consider the two following conditions: (a) You are told in advance the green-to-blue ratio in the box, e.g., that it is “50–50”; or (b) You are not told the green-to-blue ratio. Your physical activity is identical in the two cases—scooping and then recording the outcomes. Moreover, if you were to plot your outcomes, say by the number of green marbles in the sample, you would receive the same outcome distribution under either condition. Yet, could it be that in the former case you are performing a probability experiment, whereas in the second case you are engaging in statistical activity? Furthermore, could it be that one and the same sample-mean distribution demonstrates the Law of Large Numbers under the ‘probability’ condition—the cumulative distribution progressively converges on the binomial—, whereas the distribution demonstrates the Central Limit Theorem under the ‘statistics’ condition—a statistical value that is non-uniformly distributed in a population “becomes” bell shaped? If this exercise is

of any inherent value; if it enfolds the potential for deep understanding of randomness, sampling, and distribution, what media could best enable students to think about these issues?¹

Objective

That form follows function and function follows form in the evolution of human artifacts; that the message is inextricably bound to the medium in human communication and creativity; and that content and technology co-evolve iteratively in tight reciprocal feedback loops have all been demonstrated in historical analyses of human literacy (Olson, 1994), art (Jay, 1988), science (Benjamin, 1968; Gigerenzer, 1994; Latour, 1990; Wolfram, 2002), media studies (McLuhan, 1964), and numerical inscription systems, the latter through archeological (Ifrah, 2000; Schmandt-Besserat, 1992) as well as ethnomethodological studies (Saxe, 1981; Saxe & Esmonde, in press). Informed by these dynamic form–content reciprocities, education visionaries have called to align content, mathematical and scientific inquiry practices, and curricula vis-à-vis the recent computation revolution (diSessa, 2000; Papert, 1980; Wilensky, 1993, 2006; Wilensky & Papert, 2006; Wilensky et al., 2005). The argument is that canonical constructs can become imbued with new meaning—familiar phenomena can be *restructured* in radically different epistemological frameworks—when constructed within computational media (Wilensky, 2006; Wilensky & Papert, 2006; Wilensky et al., 2005). For example, the geometrical concept of circle can be construed as an embodied and dynamic construction procedure, e.g., “repeat 360 [forward 1, left 1],” rather than as a set of points equidistant from a center, when the circle is experienced through the *Logo* turtle (Abelson & diSessa, 1986; Papert, 1980). It might look like a circle and smell like a circle, but it’s a new kind of circle.

This paper looks at the potential impact of a particular type of computational applications, multi-agent modeling-and-simulation environments, on learners’ understanding of particular mathematical domains, probability and statistics. Specifically, we will examine how aspects of a computational environment designed for learning both probability and statistics—the environment’s interface features, underlying procedures, and interaction affordances—restructure the relation between probability and statistics in ways that challenge historical distinctions between these domains of study and practice. That is, unlike traditional media, the computational medium enables an equivalence class in which “strictly probability” or “strictly statistics” activities can be told apart only by virtue of disambiguating their activity contexts. For example, a user may interactively generate a data set that could be construed either as an outcome output by the simulated probability experiment or as a statistical sample from a population. I will examine what it is about the nature of the computational environment that enables probability and statistics to trespass each other’s ground. Furthermore, whereas the probability–statistics activity-based homology was serendipitous—it emerged through my design-research work within the affordances and constraints of the environment—I will argue that this homology need not be regarded as potentially confusing and even detrimental to learning but rather that it may hone learners’ understanding of fundamental ideas of these

¹ This paper emanates from an on-going design-based research on computer-based interactive mathematics learning environments. Design-based research is an on-going process with no definitive ‘end.’ Some of the design ideas I discuss have been implemented, some are under development, some are planned, and some are being considered. In this manuscript I engage in ‘humble theory’ (Cobb, Confrey, diSessa, Lehrer, & Schauble, 2003) that reflects my learning about content, my learning about technology, and my learning about learning.

domains. From that perspective, I will suggest that the computational medium affords a campground and *lingua franca* wherein researchers and practitioners of traditionally distinct mathematical and scientific practices can dwell and explore deep structural and functional affinities (Abrahamson & Wilensky, 2005b; Goldstone & Wilensky, under review).

Agent-Based Modeling

This paper develops pedagogical and cognitive arguments for the potential value of an agent-based perspective on mathematics. Agent-based modeling environments (e.g., *Swarm*, Langton & Burkhardt, 1997; *Repast*, Collier & Sallach, 2001; *NetLogo*, Wilensky, 1999a) are often associated with research on scientific phenomena. Perhaps less familiar is the application of these perspectives and methodologies in mathematics research and, specifically, in the simulation of mathematical phenomena in the context of mathematics-education research (e.g., Wilensky, 1993; Wilensky, 1995, 1997; Wolfram, 2002), e.g., probability-and-statistics phenomena. Unlike in the simulation of scientific phenomena, where randomness may be essential yet backgrounded as an auxiliary factor governing the assignment of values to agent variables, the domains of probability and statistics foreground randomness, sampling, and distribution per se as target elements of study. Wilensky (1997) has argued for the cognitive purchase on probability inherent in conceptualizing frequency distributions as emerging from myriad local interactions of rule-governed agents. This paper goes further to submit that an agent-based perspective on distribution enables insight into the “nuts and bolts” of probability and statistics—insight that restructures central constructs in these traditionally related domains in deep, illuminating, and integrating ways.

As a case study, I examine connections between the Law of Large Numbers and the Central Limit Theorem and focus on the ambiguity of the construct “population” as constructed through interaction with computer-based simulations. I argue that: (a) a collection of x independent outcomes generated by a stochastic device can be interpreted either as a sample of size x out of a “population” or as a single compound event with a property indexed by the sample mean; (b) appreciating this ambiguity is difficult without computational environments; and (c) this ambiguity, though perhaps initially confusing, may potentially foster a more nuanced, connected, and generative understanding of the domain—an understanding that can then be engaged outside of the computational media, such as in the case of the marble box that opens this paper.

Design

This paper springs from two sources: pragmatics of design (ProbLab, Abrahamson & Wilensky, 2002) and structuration theory (Wilensky, 2006; Wilensky & Papert, 2006). The current section overviews the design of ProbLab and then focuses on two simulations, *Stochastic Patchwork* (probability) and *S.A.M.P.L.E.R.* (statistics).² But first, I explain how working with multi-agent modeling-and-simulation environments fosters opportunities for restructuring traditional content. At least, how such work fostered my own learning.

² To download these and many more models, go to <http://ccl.northwestern.edu/netlogo/>. To interact with online applet of the ProbLab simulations, go to <http://ccl.northwestern.edu/curriculum/ProbLab/>

Creating ProbLab: How I Came to Examine the Probability/Statistics Overlap

Overview. The perceptual and procedural similarity of probability and statistics activities, which underlie the conceptual similarity discussed in this paper, emerged somewhat fortuitously, as I explain in this section.

I have been engaged for several years in design research on computer-based simulations of probability-and-statistics activities in *NetLogo* (Wilensky, 1999), a multi-agent modeling-and-simulation environment. These activities were created under the umbrella of the *Connected Probability* project (Wilensky, 1997) and implemented in middle-school classrooms (Abrahamson, in press-b, in press-c; Abrahamson, Janusz, & Wilensky, in press; Abrahamson & Wilensky, 2003, 2004a, 2004b, 2004c, 2005a, 2005c, 2005d, 2005e, in press).

Generic visuals and learning. The interface objects in the ProbLab simulations—the stochastic devices (probability) and the populations (statistics)—are generic yet aesthetic, e.g., a virtual mosaic of squares that can each be either white or black. My decision to work with generic materials rather than rich “situated” objects was informed by a design philosophy that “neutral” materials are more generative in terms of: (a) eliciting a diversity of students’ personal interpretations; (b) suggesting relations to other mathematical concepts; and (c) enhancing the prospects that students will correctly model future encounters with suitable situational contexts (Abrahamson, in press-a; Uttal & DeLoache, 1997). The model is designed to strike a generative balance, a middle ground, between over-particularity of cases and potential opaqueness of mathematical inscription.

Working with generic forms both in the probability and statistics activities increased the likelihood that the same forms would be used for both activities, and this perceptual overlap suggested that deeper connections should be explored. Elsewhere (Abrahamson, in press-c), I refer to the interface objects as *computational pictographs* and discuss how these objects can be designed either to foreground or background relations between mathematical constructs, such as combinatorial space and outcome distributions.

Programming and learning. In addition to the perceptual similarity of the probability-and-statistics interface objects, there is strong procedural similarity in the modeling “code” underlying and mobilizing interface elements of these simulations. The medium-specific constraint of computational modeling environments—that interface activities have to be articulated as programmable procedures—is conducive to mathematical and scientific insight (Papert, 1980, 1991, 1996). Indeed, the ‘primitives,’ commands, and procedures (the “semantics, syntax, and narrative”) that I used for the probability-and-statistics simulations have much in common. This procedural similarity is not an artifact of the perceptual similarity of the interface objects, because the procedures enfold the *behavior* of the interface entities (the “agents”) and not just their appearance.

On media, implicit constraints, and insight. At this point (see next section) it will be helpful to describe two typical ProbLab activities—one in probability and one in statistics—so as to demonstrate the perceptual and procedural similarity underlying the conceptual connections between the content as it is mediated through these activities. In so doing, I will examine the work of the designer, who builds the model through iterative authoring and interacting, as well as the work of the “end user,” who typically works within a more constrained parameter space (although she can always delve into the underlying procedures and modify them). Yet to begin with, it is important to note that both the probability and the statistics simulations are lodged in the same modeling-and-simulation environment, *NetLogo* (Wilensky, 1999).

One might gloss over the shared infrastructure of the probability and statistics activities as perfunctory; as though a computer-based environment is “just another medium”; as though “what matters is what happens *within* this medium.” However, the implicit constraints a designer confronts in developing different activities within a single medium create an a priori common design language that is informed by the medium’s logical structure. By virtue of working with this given language—the textuality of the interface visual semantics and underlying in-built syntax (Melville & Readings, 1995)—insights may emerge into potential structural relations between the content, much as working within the medium of origami may suggest topological affinities between a hat and a crane (whether the bird or the construction machine). My argument is that a computer-based medium for reasoning through expressive construction, as compared to the paper-and-pencil medium, has unique characteristics in terms of supporting the detection of conceptual affinities between different mathematical activities. These characteristics of the computation-based learning environment play out differently as contingent on the agency of the user, which may range from building the model from scratch to using a ready-made model, as I explain below.

Just as a traditional mathematician uses mathematical inscriptions (“math”) to think through problems, a modeling mathematician uses modeling procedures (“code”) to construct a simulation. These persons may be working within the same content domain, yet whereas the traditional mathematician is using a self-referring semiosis to manipulate lemmas toward examining conjectured theorems (Rotman, 2000), the modeling mathematician is creating a dynamic artifact that reconstructs selected aspects of practices and objects *using the visual expression of the practices and objects themselves*. Thus, the modeling environment affords opportunities to bootstrap mathematical knowledge through wrestling with embedded constraints of familiar objects and phenomena—the modeler need not initially be versed in the mathematical content he is learning, because the content will emerge as a problem-solving tool the modeler appropriates in improving a fit between the simulation (what occurs on the interface) and the simulated (the source phenomenon being modeled) (Papert, 1980; von Glasersfeld, 1987). This picture is somewhat different for an end-user who is not active in constructing the learning environment, as follows.

Granted, within a paper-and-pencil medium one could inscribe mathematical expressions that inform insight into the relation of different phenomena, e.g., probability and statistics, yet these expressions are a medium within a medium—mathematics within paper-and-pencil—and an end-user would need to come into the learning environment already sufficiently versed in mathematics in order to appreciate the affinities between the domains. A computer-based modeling-and-simulation learning environment, e.g., NetLogo, invisibly scaffolds potential insight into conceptual affinity by creating common visual objects. Moreover, the simulations can be calibrated so as to strike a pedagogically informed balance between “giving it away” and “discovery learning.” For example, the visual resemblance between probability-related and statistics-related interface objects could be tuned so as to control the extent to which attending to surface features may suggest their deeper conceptual connections. Moreover, just by virtue of grouping probability and statistics activities, the designer sends the user an implicit didactic directive.

Whereas all probability and statistics simulations in ProbLab share a great deal of underlying procedures (the ‘code’), these simulations do not all share interface features, because the contexts they simulate diverge. In what follows, I have intentionally selected a probability simulation and a statistics simulation that share much perceptual (interface) features. I have

selected thus to support the rhetorical thrust of this paper. Finally, note that the objective of this paper is primarily to discuss the restructuring of probability and statistics. For further details on the design, see Abrahamson and Wilensky (2004c; , 2006) and Abrahamson, Janusz, and Wilensky (in press).

Probability: Stochastic Patchwork and the Law of Large Numbers

We will begin with a definition of probability and then discuss the design of the ProbLab model *Stochastic Patchwork*, focusing on the representation of the Law of Large Numbers.

“Probability is the branch of mathematics that studies the possible outcomes of given events together with the outcomes’ relative likelihoods and distributions. In common usage, the word “probability” is used to mean the chance that a particular event (or set of events) will occur expressed on a linear scale from 0 (impossibility) to 1 (certainty), also expressed as a percentage between 0 and 100%. The analysis of events governed by probability is called statistics” (Weisstein, 2006b).

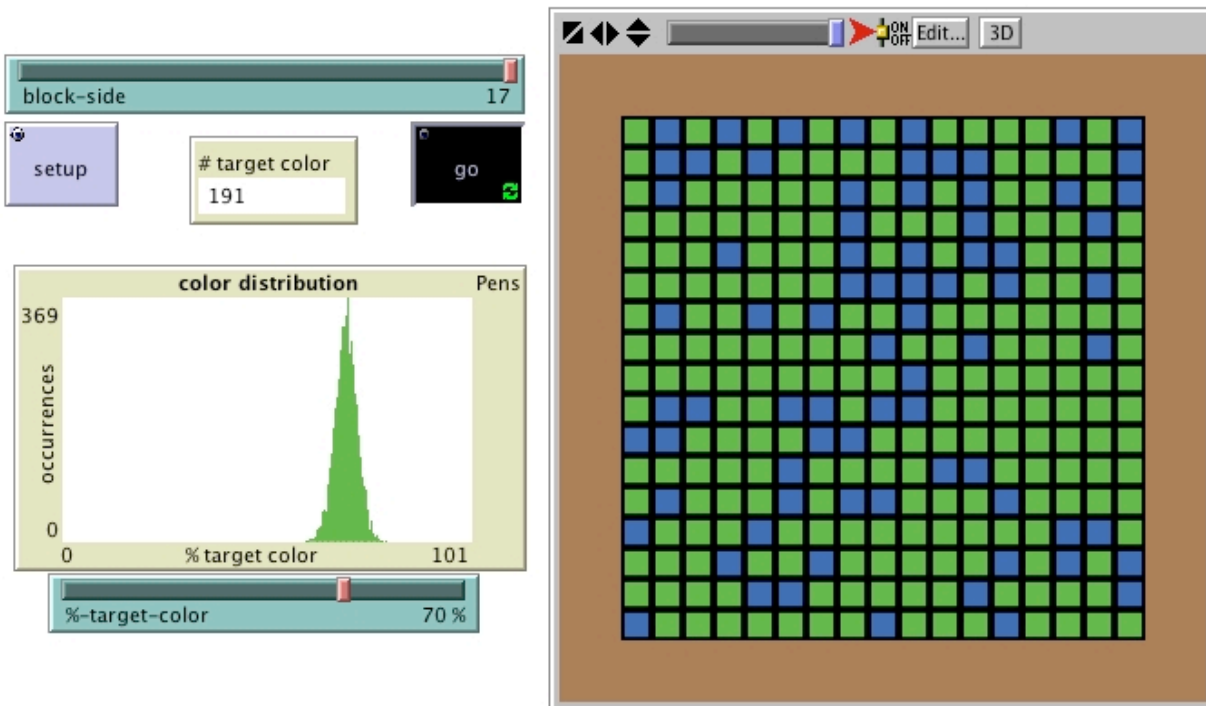


Figure 1. ProbLab: Stochastic Patchwork. Parameters are set so each square in the “mosaic” (total of $17^2 = 289$ squares) has an independent .7 (or 70%) chance of being green on each trial. Therefore, on each trial, the mosaic has approximately .7 green squares, and after several thousand trials the histogram approximates a normal distribution converging on .7.

Overview of the Stochastic Patchwork model. *Stochastic Patchwork* (see Figure 1) is an interactive computer-based model for the simulation of probability experiments. Following a brief overview, below, I will explain the objectives of the design. The randomness device—the equivalent of a coin, die, or spinner—is an array of squares that can each be either green or blue (so each square is a coin of sorts). In Figure 1 there are 289 squares, and at this “moment” 191 of them chanced to be green (green is the “target color,” i.e., the favored event monitored in the

experiment). Typical of NetLogo models, the Stochastic Patchwork interface is fully modifiable and features, amongst other tools:

- (a) a *view* window (where the stochastic array is displayed);
- (b) *buttons* for executing commands (Setup and Go)
- (c) *sliders* for setting values along core dimensions: the number of events in the array (e.g., 289, the square of ‘block-side’ 17, in Figure 1) and the p value that governs the independent chances for each cell to be the favored event (being green; set to 70% or .7 in the figure);
- (d) a *monitor* for displaying real-time and/or cumulative results (e.g., “# target color,” in Figure 1, shows “191,” the number of squares in the stochastic array that are currently green);
- (e) a *plot* window representing cumulative results (the “color of distribution” histogram shows the distribution of the percentage of green squares within the array over successive trials; 191/289, the current value, is $\sim .66$, which is just under the expected value of .7; thus, briefly before the picture in Figure 1 was taken, the histogram column directly above the abscissa value “66” grew vertically by one unit);

Not seen here are the procedures governing the behavior of the model (they are in a separate tab). These procedures are fully editable and extendable, thus enabling further parameterization of the experiment and other modification emerging throughout the inquiry process. I now further explain how a user might construe the simulation.

The fundamental probabilistic element in this model is a square “coin” that has a green “side” and a blue “side.” Instead of flipping this single coin many times, we flip many clones of this coin all at once. The crux of this model is that if a single coin has a .7 chance of landing on green, then the aggregate of a sufficiently-large sample of these coins that flip all at once will approximate a .7 greenness, i.e. most of the time about .7 of the squares will be green. Because the Stochastic Patchwork model simulates a probability experiment, outcomes vary, yet after a sufficiently large number of successive iterations in the experiment the outcome *distribution* approximates a normal curve converging on a .7 mean value of greenness, as displayed in a histogram that is part of the model (see Figure 1). The objective of students’ interacting with the model is that they understand how the model works, possibly through exploring the effect of modifying parameter settings—the size of the array and/or the bias of the coin/square—on the dynamics of the emerging distributions.

We have found that students as young as 10-years old working with ProbLab models interpreted experimental results using both enumeration-based strategies (counting green and blue squares in samples) and multiplicative strategies (inferring proportions in populations by eyeballing green/blue ratios in samples)(Abrahamson, Janusz, & Wilensky, in press).

Having described the basic apparatus, I will now focus on the representation of the Law of Large Numbers in this simulation. First, a definition and then an intuitive explanation that grounds the theorem in the Stochastic Patchwork model.

“A “law of large numbers” is one of several theorems expressing the idea that as the number of trials of a random process increases, the percentage difference between the expected and actual values goes to zero” (Weisstein, 2006a).

A qualitative explanatory model for intuitively grasping the Law of Large Numbers might focus on an interpretation of the outcome distribution as a transformation on the combinatorial space (see below). Note that this explanation assumes some familiarity with combinatorics

Stochastic Patchwork and the Law of Large Numbers. There are 2^{289} unique configurations of a stochastic array of 289 cells each of which can be either green or blue. The number 2^{289} is of the order of magnitude of 10^{87} , that is, “a 1 with 87 zeros to its right.” Each of these numerous unique configurations has the same likelihood— $1/2^{289}$ —to occur at each independent experimental trial. Yet the focal statistical value in the experiment is not *specific* configurations but the total number of green cells in each of these configurations. That is, in mathematizing the array outcomes, we collapse the spatial element, reducing the visual display to a single value, e.g., 191 (see Figure 1). However, the spatiality of the array is key to understanding the emergent outcome distribution. That is, the statistical analysis may ignore the spatial dimension, yet it is this dimension that underlies the focal concept embedded in the simulation, namely the Law of Large Numbers. Thus, the array is not a mere “visualization” of some deeper or truer mathematical phenomenon—it is the phenomenon itself that is being analyzed.

Out of all the 2^{289} unique configurations, there is only one way of “getting” an all-blue array (with no green cells). Yet to get an array with a *single* green square there are 289 different configurations, because each one of the 289 cells might be green while the other 288 cells are blue. For *two* green squares there are $289 \times 288 / 2$, or 41,616 different arrangements. So we have gone up in the sequence 1, 289, 41616.... Over a stupendous number of trials, each one of the 2^{289} unique configurations appears roughly as many times as the others—at least, the expected distribution by *unique* configuration is flat. Let us think of this flat distribution as comprised of 2^{289} “units.” The Stochastic Patchwork histogram can be construed as a manipulation of these units. We stack these numerous units according to our statistic, the number of green cells, resulting in only 290 columns (because there can be 0 green cells, 1 green cell, 2, 3...288, 289). In so doing, for some of the 290 columns we will harvest more units from the flat distribution than for other columns. For example, for the no-green column we will harvest just one unit, but for the one-green column we will harvest 289 units. Thus, within this new distribution, the more there are unique possible configurations in a class of events, the taller the pooled histogram column becomes. For example, if we pool the occurrences of each and every one of the 289 configurations with exactly one green cell in them, we will get a histogram column that is roughly 289 times as tall as the no-green column—it has 289 units as compared to a single unit. This idea has been further explored in interview-based studies with elementary- and middle-school students (Abrahamson, 2006), using a 2-by-2 array (the *4-Block*) that has a total of only 16 unique possible outcomes.

Statistics: S.A.M.P.L.E.R. and the Central Limit Theorem

We will begin with a definition of statistics and then discuss the design of the ProbLab model *S.A.M.P.L.E.R.*, focusing on the representation of the Central Limit Theorem.

Statistics is, “The mathematical study of the likelihood and probability of events occurring based on known information and inferred by taking a limited number of samples” (Weisstein, 2006c). Note that, unlike the case of probability experiments where the experimenter may know the properties of the stochastic device, in statistical practice the experimenter infers the population’s properties from the samples. We will return to this point in the next section.

Overview of the S.A.M.P.L.E.R. model. *S.A.M.P.L.E.R.*, Statistics As Multi-Participant Learning-Environment Resource, is a *participatory simulation activity* (Wilensky & Stroup, 1999b) implemented in the *HubNet* (Wilensky & Stroup, 1999a) technological infrastructure,

which extends NetLogo so as to enable facilitation of collaborative inquiry in a networked classroom.

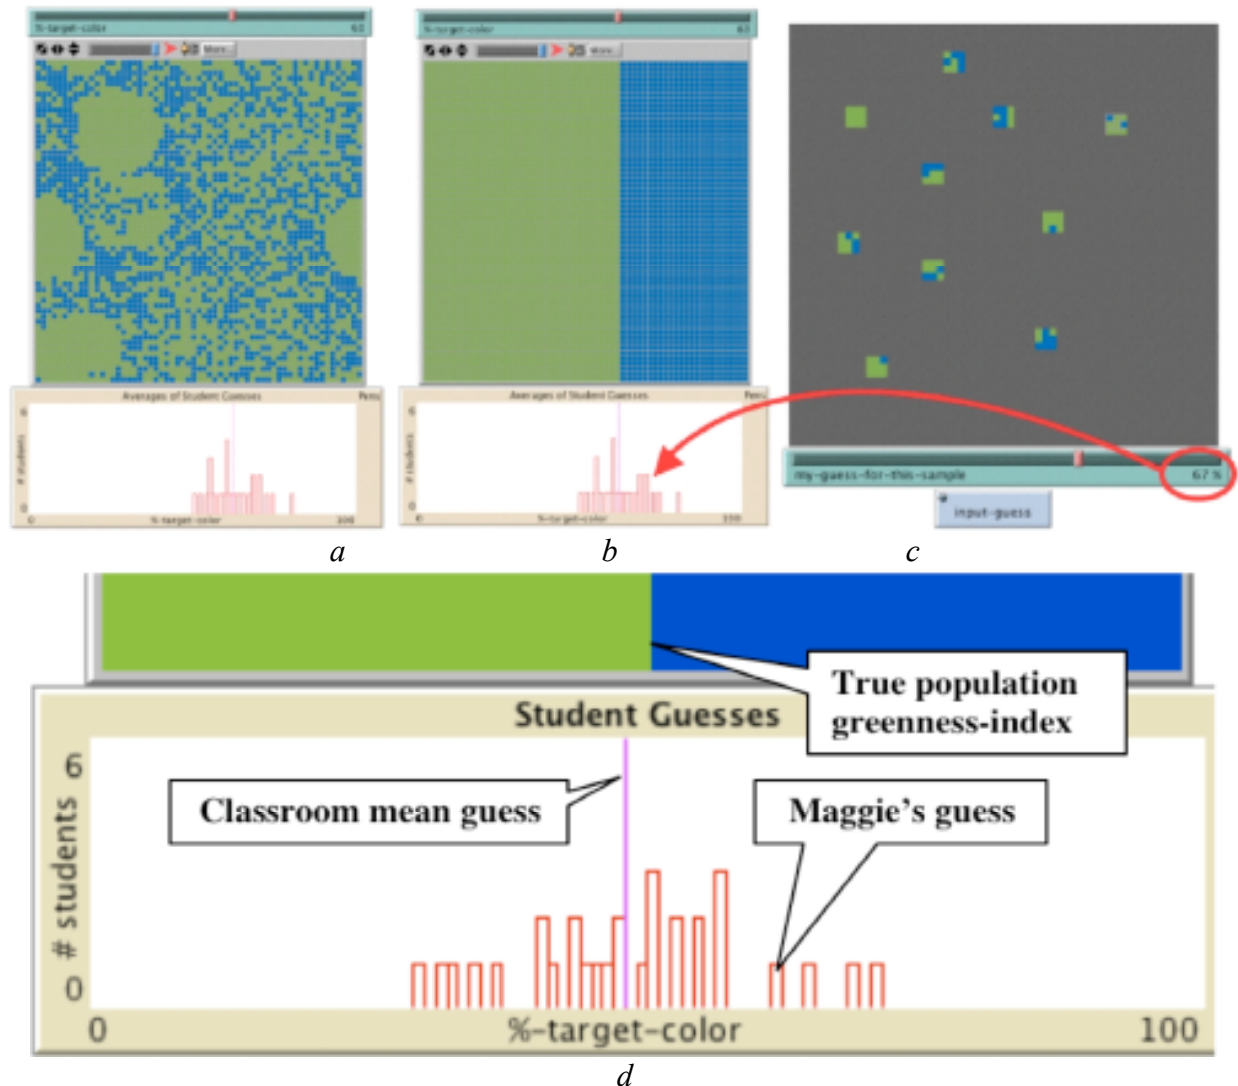


Figure 2. Selected features of the S.A.M.P.L.E.R. computer-based learning environment.

In S.A.M.P.L.E.R. (see Figure 2, above), students take individual samples from a population so as to determine a target property of this population. The “population” is a matrix of thousands of green or blue squares (Figure 2a) and the target property being measured is the population’s greenness, i.e., the proportion of green in the population. A feature of the activity is that population squares can be “organized”—all green to the left, all blue to the right (Figure 2b). This “organizing” indexes the proportion of greenness as a part-to-whole linear extension that maps onto scales that are both in a slider (above the population) and in a histogram of students’ collective guesses (below the population). The population can be set to bear an unknown random percentage of green squares, and the population can then be hidden (masked) so that information about the green/blue properties of the squares can be gleaned only through sampling. Students participate through *clients* (students’ calculators or personal computers). These clients are

hooked up to the facilitator's *server*. Students take individual samples from the population (Figure 2c), and analyze these samples so as to establish a best guess for the population's target property, its greenness. Note that whereas all students sample from the same population, by default students each see only their own samples, unless these are "pooled" on the server. Students input their individual guesses, and these guesses are processed through the central server and displayed as a histogram on the server's interface that is projected upon the classroom overhead screen (see fragment of this projection in Figure 2d).

The projected histogram shows all student guesses and the classroom mean guess, and this histogram interfaces with the self-indexing green-blue population. Note the small gap (Figure 2d, middle) between the classroom mean guess and the true population index. Because a classroom-full of students takes different samples from the same population, the histogram of collective student input typically approximates a normal distribution and the mean approximates the true value of the target property being measured. The students themselves constitute data points on the plot ("I am the 37"... "So am I!"... "Oh no... who is the 81?!"). So students can reflect both on their individual guesses as compared to their classmates' guesses and on the classroom guess as compared to the population's true value of greenness. Such reflection and the discussion it stimulates are designed to foster opportunities for discussing and understanding typical distributions of sample means (for further description of the design and of data from implementations of the design in middle-school classroom, see Abrahamson & Wilensky, 2004c). These typical distributions are predicted by the Central Limit Theorem.

The Central Limit Theorem is:

a statement about the characteristics of the sampling distribution of means of random samples from a given population. That is, it describes the characteristics of the distribution of values we would obtain if we were able to draw an infinite number of random samples of a given size from a given population and we calculated the mean of each sample (The_Animated_Software_Company, 2006).³

In the next section, we will return to the Central Limit Theorem and compare it to the Law of Large Numbers.

Having introduced both Stochastic Patchwork, a probability activity, and S.A.M.P.L.E.R., a statistics activity, we will now turn to a comparison of these designs along the dimensions of interface functionalities, student activities, and modeling code. The comparison will prepare an explanation of "bottom-up stats" and a set of examples as well as implications for design.

³ The citation continues thus: "The Central Limit Theorem consists of three statements:

1. The mean of the sampling distribution of means is equal to the mean of the population from which the samples were drawn
2. The variance of the sampling distribution of means is equal to the variance of the population from which the samples were drawn divided by the size of the samples.
3. If the original population is distributed normally (i.e. it is bell shaped), the sampling distribution of means will also be normal. If the original population is not normally distributed, the sampling distribution of means will increasingly approximate a normal distribution as sample size increases. (i.e. when increasingly large samples are drawn)."

Bottom-Up Stats

In this paper, we are examining how the practice of agent-based modeling may inform the modeler—whether a student, a designer, or researcher—of mathematical connections that are not typically salient when working within traditional media and contexts. This section examines the specific attributes of the NetLogo learning environment and ProbLab simulations that enable insight into the relations of probability and statistics. From this perspective, we then compare between the Law of Large Numbers and the Central Limit Theorem so as to prepare a discussion of probability as “Platonic statistics.” The section ends with a demonstration of how normal distributions could be construed as special cases of the Law of Large Numbers.

Comparison of the Probability and Statistics Designs: Interface, Activities, and Modeling Code

In the previous sections, I explained the Stochastic Patchwork and S.A.M.P.L.E.R. models separately. I will now compare these activities along the dimensions of their respective interface, underlying modeling code, and the users’ actions.

Interface. In both Stochastic Patchwork and in S.A.M.P.L.E.R., the smallest mathematically meaningful unit—the generic computational pictograph of these designs—is the NetLogo *patch*. The patch, an autonomous square area of interface real estate, is an “agent”—it has its own properties that are implemented as variables with designated values, such as “patch-color = green.” The patch can assign itself random values, too—it can “flip a coin”—e.g., the patch can “choose” randomly between 0 and 1 and then, based on the outcome, assign itself the color blue or green, respectively. To simulate, within an *electronic* medium, the generation of a random outcome with a *physical* stochastic device, the patch activates a computer-based randomness emulator routine.

Patches are immobile agents, unlike the “turtle,” a mobile agent (Wilensky, 2001). In the probability-and-statistics models discussed here, I have used only patches (and not turtles). In other ProbLab models, turtles, too, carry random values, and sometimes these are combined with patches, e.g., a turtle will “carry” a randomly-assigned value to a patch. Elsewhere (Abrahamson, in press-c), I discuss the potential of immobile computational pictographs as perceptual–conceptual bridges from raw data to mathematical representations, i.e., the data are stacked up in columns to “become” the histogram. Also, note in Figures 1 and 2a how the green and blue patches cover the totality of the view window, i.e. every single patch in that “world” is either green or blue. Gazing at this spatial representation either as at an outcome distribution (probability) or as at a population (statistics), one can engage perceptual-based proportional reasoning to establish a qualitative sense of the color ratios, either as part-to-part or as part-to-whole (Abrahamson & Wilensky, 2004b). When mobile agents (turtles) are used, a third color—the necessary background color—may hamper such perceptual judgment.

In the process of designing ProbLab, the convergence of both the probability and statistics interfaces on the same object, the bi-colored patch, as the generic simulation agent, was not premeditated. Rather, this was an emergent result of working on both simulations within the same medium, with its inherent affordances and constraints, and of exploring within each simulation the prospects of engaging users’ perceptual judgments as well as connections to standard mathematical representations.

Modeling code. Table 1, below, compares two modeling procedures.⁴ Even a person with little or no experience in authoring within programmable environments will notice a resemblance between these juxtaposed procedures. However, these procedures appear at different points within the enactment of the code, and this difference is important for our discussion, as follows.

Table 1.

Comparison of Probability and Statistics NetLogo Simulations by Modeling Procedures That Determine Outcome Values

Stochastic Patchwork (Probability)	S.A.M.P.L.E.R. (Statistics)
to go ;; observer procedure	to initialize-patch-colors ;; patch procedure
ask patches-in-block [
ifelse random-float 100 < %-target-color	ifelse random-float 100 < real-%-target-color
[set pcolor target-color]	[set my-color target-color]
[set pcolor other-color]	[set my-color other-color]
]	
end	end

The Stochastic Patchwork procedure, “to go,” embodies instructions to the cells: Each cell selects a random number between 0 and 99; and if this number is lesser than some designated value that is set on the interface, then the cell becomes green (the target color)—otherwise, it becomes blue (the other color). This procedure embodies the act of generating a random compound-event sample (“rolling” or “flipping” the array)—it is the “go” procedure, i.e. the real-time “run” through the simulation. The S.A.M.P.L.E.R. procedure is essentially identical, yet this procedure is an initializing procedure, i.e., it is run prior to the “action” so as to set up the backdrop for the simulation. That is, whereas in Stochastic Patchwork there is only a direct enactment of a random compound event, in S.A.M.P.L.E.R. the sampling activity is decoupled into an initializing stage where the population is created and a sampling stage, where selected parts of this population are exposed. Between these two stages, the population is “hidden,” i.e., the cells don a masking color and will don their real color—the ‘my-color’ variable in Table 1—only if the user selects them.

Activities. In both the probability and statistics simulations the user’s basic action—whether it is pressing the ‘Go’ button (probability) or clicking on a blank space (statistics)—causes the patches to reveal their current color. Moreover, in both simulations the samples are taken successively, constructing a sequence of data that is translated into a mathematical representation (the histogram). Thus, the spatial array is generated or parsed in a temporal sequence—trial by trial or sample by sample—which is then deployed into a different cumulative spatial form. That is, *the temporal medium enables a bridging between two spatial representations—the generic form (the raw data) and the mathematical form (the histogram).* Furthermore, just as statistics governs the analysis of probability experiments, so probability governs the selection of statistical data.

We have now compared the probability and statistics activities along the dimensions of interface design, programming, and user actions. Below, I use this similarity to illuminate how engaging in the probability-and-statistics computer-based activities may support a honing of

⁴ I have removed code lines that are irrelevant to this discussion, yet the reader can access these files at <http://ccl.northwestern.edu/netlogo/>.

these domains' central constructs. Specifically, Table 2 (see below) summarizes a comparison between the Law of Large Numbers and the Central Limit Theorem from the perspective of the activities we have been examining.

Table 2

A Comparison Between the Law of Large Numbers and the Central Limit Theorem From the Perspective of the ProbLab Simulations

Construct	Law of Large Numbers	Central Limit Theorem
Gist	When probability experiments are run trial by trial, the outcome distribution progressively converges on the distribution anticipated by combinatorial analysis. That is, even as the <i>number</i> of occurrences diverges from the expected values, the <i>frequencies</i> (proportional) converge.	Regardless of the distribution of a population across some finite range of values, an accumulation of sample means (or sums) from this population progressively converges on a normal distribution. This sample-value distribution has the same mean as the population and a variance that is smaller by a factor of the sample size.
Sample	An outcome generated on-the-fly by a (simulated) stochastic device. The device may produce a compound event consisting of many independent events. Metaphorically speaking, sampling is an act of "selecting" out of the combinatorial space, yet the outcomes do not exist physically prior to the sampling act but only as a potentiality, an eventuality.	A random selection of items from a population. The sampled specimens and, in fact, the entire population, exists prior to the sampling act. Not all specimens are sampled—just a fraction that is judged as large enough to inform the statistician with sufficient confidence of the target values in the population.
Resampling	The same stochastic device is used repeatedly—changing the device would undermine the logic of the experiment.	Population specimens can, in principle, be resampled, yet that would waste resources without increasing statistical power.
Distribution	Converges on 'normal,' with a smaller variance for a greater number of independent outcomes in the compound event.	Converges on 'normal' (bell shaped) with a smaller variance for larger sample sizes.
Randomness	The likelihood of events in the combinatorial space is governed by specifiable probabilities that amount to 1.0. These likelihoods could be equally distributed between all possible events, e.g., in a "fair coin." Over time, the outcomes occur stochastically.	All specimens within the population have equal likelihood to be sampled. In computer simulations, randomness may be at play twice—both in creating the population (automatic procedure) and again in selecting specimens (user action).
Activity	Mostly experimentation (albeit one could explore a stochastic device of unknown probabilities, e.g., an unfair coin)	Mostly exploration (EDA; albeit one could, in principle, experiment with drawing samples from a population with known statistics).

I have compared between the Law of Large Numbers and the Central Limit Theorem from the perspective of the corresponding computer-based activities in ProbLab, so as to demonstrate how the medium foregrounds conceptual affinities between these domains of mathematical practice. I now turn to discussing how this comparison may inform design that capitalizes on the concepts' underlying similarities so as to structure probability as a form of "Platonic statistics."

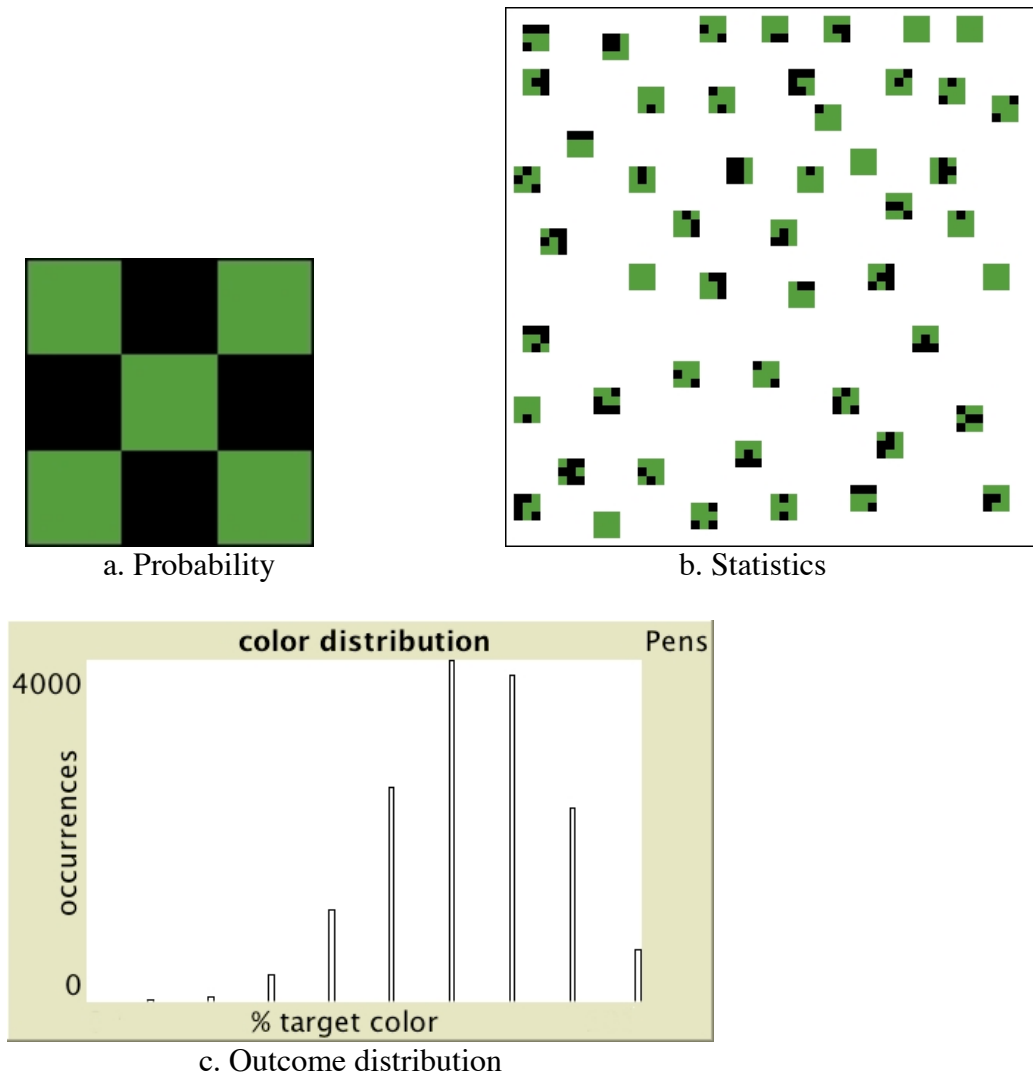


Figure 3. Probability as "Platonic statistics." *a.* The Stochastic Patchwork device calibrated at size 3-by-3 to generate a compound event with 9 independent outcomes (each cell can be either light or dark at a .7 chance)—these samples are not determinable prior to the sampling action. *b.* The S.A.M.P.L.E.R. interface with a .7 green population in which fifty 3-by-3 samples have been exposed. The exposed cells "carried" their respective color values prior to the sampling act. The probability experiment (*a.*) and the statistical sampling (*b.*) result in an identical distribution (*c.*), demonstrating the relationship between the Law of Large Numbers and the Central Limit Theorem. One could imagine a space with all the 512 possible 3-by-3 blocks—generating a random outcome with the device (*a.*) would be analogous to selecting randomly from this field.

The “Platonic Combinatorial Space” as a Bridge Between Probability and Statistics

When a stochastic device generates an outcome, the activity could be construed as randomly selecting a sample from a Platonic combinatorial space (see Figure 3, above). To use a more familiar stochastic device, flipping a fair coin could be construed as choosing from the “Heads; Tails” space (a space containing two individual coins, one showing Heads, and one showing Tails). Flipping two coins would be likened to selecting from a Platonic space containing four unique pairs: “Heads, Heads; Heads, Tails; Tails; Heads; Tails, Tails.”

I should emphasize that I am using the term ‘Platonic’ as an interim rhetorical construct. This is a didactic and not a philosophical position—I would not want to be interpreted as ascribing objective realness to such a space but rather to explore possible benefits, to learning, of reifying the potentiality of random selection; of treating the combinatorial space as a population. Perhaps more terrestrial an interpretation than a Platonic space would be that flipping the coins is analogous to drawing the outcome out of a hat containing one of each of the possible compound events (Abrahamson & Wilensky, 2005e). Yet this Platonic space or hat is homologous with the device's combinatorial space only for the special case where all outcomes are equally likely. Flipping an “unfair” coin with a .75 chance of getting Heads would be analogous to choosing from a “corrected” space, “Heads; Heads; Heads; Tails.” When one constructs (mentally, physically, or virtually) the combinatorial space as an actual population that is “out there,” probability experiments become “Platonic statistics.”

To the extent that this mental construction of probability—selecting randomly from a Platonic combinatorial space—is one that is conducive to learning the content of probability, the question then arises as to media, designs, and activities that can mediate this construction. It is here that computer technology can be uniquely beneficial. For example, in the absence of computational media, it could be very cumbersome to generate and manipulate the corrected space for flipping 2 coins where each coin has a .6 chance of landing on Heads, where the constraint is that only whole coins can be used—fractions of coins would not be accepted, so the lowest integer expression would be 4:6:6:9 for the four possible events. It is such computer-based representations, I wager, that could bridge between simple probability problems to cases with numerical values that usually require the application of advanced mathematical technology (functions)(see Gelman & Williams, 1998, on enabling constraints; see Gigerenzer, 1998, on ecological intelligence and natural frequencies). Thus, these more advanced cases could be grounded in perceptual judgment as extensions of the simpler cases. This remains to be designed and researched. For example, representing the Platonic space in the 2-coin example, above, begs the questions of whether students could follow the rationale of constructing that space and what could be done to structure it.

Media Revisited, or Once Again, Why Computers?

The insights delineated in this paper stem from several years of being immersed in the design of computer-based interactive simulations of probability- and statistics-related practice. However, it can be claimed that one could, in principle, use non-computational media to achieve similar perceptual/procedural/conceptual overlap between representations and activities for probability and statistics. Furthermore, it could be claimed that just because the designer followed a particular learning trajectory, this trajectory is not necessarily universally useful. However, outside of the computational environment such overlaps might be awkward and

ultimately confusing, because traditional stochastic (e.g., dice) and statistical (populations) objects are strongly associated with particular situational settings. For example, consider the case of dice, which are strongly associated with situations of probability. Let us attempt to construe dice as objects pertaining to a statistics activity. Imagine 625 dice arranged in a tightly knit 25-by-25 array. The array is covered by an opaque material that enables a view only of selected dice (the “sample”). One *could* construe this array as a collection of 625 specimens in a population wherein the statistic being measured is the dice’s ‘face value.’ Furthermore, one could take samples from this population to measure the distribution of face values in this population. Whereas this activity may be mathematically sensible, I wager that the activity would prove technically and psychologically cumbersome. Thus, dice’s association with probability and not statistics would impede a comparison of the concepts on neutral grounds, and similar considerations apply to many traditional stochastic generators, such as coins or spinners.

The ProbLab interfaces offer a more neutral ground. Unlike physical probability objects (e.g., dice), the virtual two-dimensional ProbLab squares do not bear all of the outcomes simultaneously as inherent properties and omnipresent potentialities. A square is either white or black. That is, the interface objects and modeling procedures hide the potentiality of these squares to bear either of the two colors, so that the interface reveals only one of these two values at any given moment, per square.

Finally, the marbles-in-a-box case that was introduced in the preface of this paper is, perhaps, an approximation of a context that enables flexibility in its framing either as ‘probability’ or ‘statistics.’ The box could be construed as a Platonic combinatorial space of a single fair coin, albeit an expansion on the 1:1 ratio of Heads to Tails, say to 100:100. Yet that case does not readily enable modification of the ratio between the color values, and this technical shortcoming delimits a user’s exploration of multiple parameter settings. The potential solution of working with small numbers of marbles is not satisfactory, because it precludes the entire statistical enterprise—if one can view the whole population in a glance, one need not take samples but just count. Such complications can be eschewed in computational environments.

Lagniappe: Constructing the Central Limit Theorem as a Case of the Law of Large Numbers

We have been discussing mathematical models that use simple visualizations and procedures—square cells (computer agents) that each “choose” to be either one color or another. I have argued that this simplicity may be generative in the sense that it may illuminate deep relations between otherwise disparate situations and mathematical ideas. Also, I have underscored the roles that technology may play in the model-based explorative inquiry processes. However, seeking deep underlying connections between concepts, as an intellectual practice, may have its limitations—it may overgenerate to the point where so many constraints must be satisfied that only obscure overlap cases are covered. And yet, it is my impression that the very practice of seeking conceptual relationships—regardless of whether meaningful relationships are found—is in and of itself conducive to the clarification of each of the compared situations; it is a form of mathematical discourse that produces well articulated definitions. I will leave it to the reader to judge whether the mathematical relation uncovered in this section, between the Law of Large Numbers and the Central Limit Theorem, is meaningful or obscure and whether, at the least, the comparison helped to clarify each of the compared concepts and to inform potentially useful design directions.

This section explores the plausibility of comparing between two different mathematical situations, one with x replications of one variable and the other with an x -variable (a compound

of x distinct types of events).⁵ The context is the bell-shaped distribution of the statistic ‘height’ within the population of 6th grade female students in the USA (we will be working with fictitious assumptions and not with real data). Instead of taking samples, say of 9 girls per sample, and plotting these samples’ means, thus converging on a bell-shaped distribution as the Central Limit Theorem predicts, we will build the distribution bottom up. Namely, we will consider each data point on the plot not as a sample of 9 girls but as a single girl whose height is determined by 9 variables. These 9 variables will be visualized using a 3-by-3 grid of cells, the *9-Block* (see previous sections).

In this basic explanatory model of height distribution in a large population, the following assumptions are made:

- Only nine variable factors, genetic and environmental, contribute to determining a person’s height, such as ‘parents’ height,’ ‘nutrition,’ ‘climate,’ etc. These nine variables map onto the nine squares in the 9-block.
- Each person is randomly assigned values for each of the nine variables.
- Instead of looking at a range of values per variable factor, we will assume that each variable can only be either a ‘positive’ (green) or a ‘negative’ (blue).
- Each variable value has an equal likelihood (.5 probability) for positive or negative, e.g., ‘nutrition’ could equally be either ‘good’ or ‘bad.’
- Each of the nine factors contributes equally towards determining a person’s height, e.g., ‘good nutrition’ would be as powerful as ‘tall parents.’
- The more positive values a person has, the taller this person is, along some interval scale.
- The variables do not interact. For example, the contribution of a positive value for ‘climate’ is not dependent or affected by the value of any of the other variable properties.
- The range of the distribution of heights is limited both on the minimum and maximum (there is a “cutoff” standard deviation from the mean).




Figure 4. The combinations tower. In the 9-Block (on the left), a 3-by-3 grid of 9 cells, each cells can be either green or blue. The combinations tower (on the right, with enlarged detail) is the combinatorial space of the 9-Block, with 512 unique configurations and arranged in columns according to the number of green squares in the blocks.

⁵ This section appears also in (Abrahamson, Janusz, & Wilensky, in press)

Given the above assumptions, the combinations tower (see Figure 4, above) is the combinatorial space of all the possible configurations of the nine variable factors contributing to a person's height, and these combinations are organized by "height groups" from 'shortest students' (on left) to 'tallest students' (on right). The leftmost 'shortest students' column and the rightmost 'tallest students' column each hold a single combination, the no-positive and the all-positive, respectively. In between, there are more combinations that yield a count of 4 or 5 positives as compared to, say, 3 or 6 positives, and there are more combinations that yield a count of 3 or 6 positives, as compared to 2 or 7, etc. Yet note that each of the 512 combinations is equally likely to occur. Thus, the bell curve can be understood as a combinatorial sample space of a cluster of variables that has been detected as contributing to a property of an observable phenomenon. Each variable is independent of the others, but as a cluster of variables that inform a property of a phenomenon, these nine variables are co-dependent, and these co-dependencies create the bell-shaped combinatorial distribution. This is why instances of a phenomenon are often distributed such that there are more "average" incidents. For example, there are more people of average height than there are short people or tall people.

By way of demonstrating how this model may be complexified, consider the dichotomous variable of each square in the generic model. If we were to increase the space of possible outcomes to 3, the combinations tower would grow to encompass $3^9 = 19,683$ different possibilities. The shape of this tower would be closer to a bell shape. If we were to modify the relative likelihoods or weight of individual squares or if we were to introduce causal contingencies between squares within the 9-block, we might affect the shape of the distribution. Implementing this model as a computer-based interactive simulation would enable us to readily explore the parameter space, tinker with the procedures underlying the emergent distribution, and receive immediate feedback in the form of mathematical representations. By way of comparing these simulated experiments to information from scientific and statistics resources, we can evaluate the explanatory power of our model and iteratively modify the model toward a better fit with the data.

Summary

Underlying the enterprise pursued in this paper is a student-centered perspective on mathematical content. It is trivial to point out that, coming into a learning environment the student does not know the content that the designer and teacher wish for her to learn. That is, the student does not have the larger-picture vantage point from which to judge what an activity is 'a case of'—the student is "doing stuff," i.e., considering situations, constructing materials, evaluating results, and discussing them. While designers must bear this in mind vigilantly, designers also have the opportunity to restructure relations between mathematical domains where such restructuring is pedagogically warranted. That is, the student-centered and not content-centered perspective on students' classroom experiences enables us to put aside considerations of targeted content and focus instead on potential connections students are discerning *across* mathematical domains, across activities—such connections may have been obfuscated by historical media but may come to light through current computational media.

In this vein, I have demonstrated how the 'student-with-computer' overlapping phenomenology in the probability simulations and statistics simulations—in terms of the actions of *obtaining samples*, in terms of the *temporal dimension* of these acts, and in terms of *visual metaphors* that appear to be apt for thinking about both probability and statistical

phenomena—restructures these formally disparate constructs as nuanced epistemological variants governed by the parameter ‘experimentation vs. exploration.’ Normatively, probability and statistics are phenomenologically distinct practices. Yet epistemologically what is at stake is not the normative activities or contexts per se but whether the person knows or does not know the central tendency values of the device or population while she is sampling. That is, a population can act as a stochastic device for probability experiments—and not as a statistical construct—as long as we know the population’s mean and distribution. Furthermore, in order to capitalize on the probability–statistics relations, I have suggested the “Platonic combinatorial space”—the fictitious population from which random outcomes are drawn in probability experiments. Finally, I presented a “bottom-up statistics” exploratory model that builds a Law-of-Large-Numbers explanation of the Central Limit Theorem.

Conclusion and Implications for Education

Radical progress in the fields of mathematics and science can be couched as restructurings of phenomena, using new insight, methodologies, and media (Goldstone & Wilensky, 2005; Kuhn, 1962; Thagard, 1992; Wilensky, 2006; Wilensky et al., 2005). In particular, current technology affords new perspectives on familiar content. From these new vantage points, conceptual distances between domains formerly regarded as disparate may prove to be not inherent to the domains but largely historical artifacts of taxonomy, methodology, and activity contexts. These conceptual distances may justly be perpetuated within the fields of *applied* mathematics and sciences, e.g., in engineering, because those fields are tied to richer activity contexts for which restructurations may not be relevant. Yet for the field of education, deeper conceptual understanding may lie in exploring the “family resemblance” of apparently disparate domains (Wilensky, 2006). Such exploration foregrounds the implicit role of historically extant media in shaping conceptual understanding; the exploration supports disengagement from the “functional fixedness” that media may instill. These days of rapid technological progress—days in which teenagers are nostalgic of audio–visual media from their own childhood epoch—are auspicious days for an educational emphasis on media–concept relations. An educational emphasis on media–concept relations may foster an epistemological disposition that is more conducive to bringing on the next conceptual revolution.

Author’s Note

This paper builds on research funded by NSF ROLE grant 0126227 (Uri Wilensky, PI). Writing this paper was enabled through the support of the National Academy of Education/Spencer Postdoctoral Fellowship 2005–6 (*Seeing Chance* project; Abrahamson).

References

- Abelson, H., & diSessa, A. A. (1986). *Turtle geometry*. Cambridge, MA: MIT Press.
- Abrahamson, D. (2006). The odds of understanding the Law of Large Numbers: A design for grounding intuitive probability in combinatorial analysis. *Manuscript under review*.
- Abrahamson, D. (in press-a). “Because in the world, there are more blocks of this type”: The real-worldness of immersive combinatorial analysis as a grounding of simulated probability experiments. In D. Abrahamson (org.), What’s a situation in situated cognition? – A constructionist critique of authentic inquiry. In S. Barab, K. Hay & D. Hickey (Eds.), *Proceedings of the Seventh International Conference of the Learning Sciences*. Bloomington, IN: ICLS.

- Abrahamson, D. (in press-b). Learning chance: Lessons from a learning-axis and bridging-tools perspective. In A. Rossman & B. Chance (Eds.), *Proceedings of the Seventh International Conference on Teaching of Statistics*. Salvador, Bahia, Brazil.
- Abrahamson, D. (in press-c). The shape of things to come: The computational pictograph as a bridge from combinatorial space to outcome distribution *International Journal of Computers for Mathematical Learning*.
- Abrahamson, D., Janusz, R., & Wilensky, U. (in press). There once was a 9-Block... -- A middle-school design for probability and statistics. *Journal of Statistics Education*.
- Abrahamson, D., & Wilensky, U. (2002). *ProbLab*. Northwestern University, Evanston, IL: The Center for Connected Learning and Computer-Based Modeling, Northwestern University. <http://ccl.northwestern.edu/curriculum/ProbLab/>.
- Abrahamson, D., & Wilensky, U. (2003). *The quest of the bell curve: A constructionist approach to learning statistics through designing computer-based probability experiments*. Paper presented at the Third Conference of the European Society for Research in Mathematics Education, Balleria, Italy.
- Abrahamson, D., & Wilensky, U. (2004a). ProbLab: A computer-supported unit in probability and statistics. In M. J. Høines & A. B. Fuglestad (Eds.), *Proceedings of the 28th Annual Meeting of the International Group for the Psychology of Mathematics Education Vol. 1* (pp. 369). Bergen, Norway: PME.
- Abrahamson, D., & Wilensky, U. (2004b). *S.A.M.P.L.E.R.: Collaborative interactive computer-based statistics learning environment*. Paper presented at the 10th International Congress on Mathematical Education, Copenhagen, Denmark.
- Abrahamson, D., & Wilensky, U. (2004c). *S.A.M.P.L.E.R.: Statistics As Multi-Participant Learning-Environment Resource*. Paper presented at the annual meeting of the American Educational Research Association, San Diego, CA.
- Abrahamson, D., & Wilensky, U. (2005a). *Collaboration and equity in classroom activities using Statistics As Multi-Participant Learning-Environment Resource (S.A.M.P.L.E.R.)*. In W. Stroup and U. Wilensky (Chairs) & C. D. Lee (Discussant), *Patterns in group learning with next-generation network technology*. Paper presented at the annual meeting of the American Educational Research Association, Montreal, Canada.
- Abrahamson, D., & Wilensky, U. (2005b). *Piaget? Vygotsky? I'm game: Agent-based modeling for psychology research*. Paper presented at the annual meeting of the Jean Piaget Society.
- Abrahamson, D., & Wilensky, U. (2005c). *ProbLab goes to school: Design, teaching, and learning of probability with multi-agent interactive computer models*. Paper presented at the Fourth Conference of the European Society for Research in Mathematics Education, San Feliu de Gixols, Spain.
- Abrahamson, D., & Wilensky, U. (2005d). *The stratified learning zone: Examining collaborative-learning design in demographically-diverse mathematics classrooms*. In D. Y. White (Chair) & E. H. Gutstein (Discussant), *Equity and diversity studies in mathematics learning and instruction*. Paper presented at the annual meeting of the American Educational Research Association, Montreal, Canada.
- Abrahamson, D., & Wilensky, U. (2005e). Understanding chance: From student voice to learning supports in a design experiment in the domain of probability. In G. M. Lloyd, M. Wilson, J. L. M. Wilkins & S. L. Behm (Eds.), *27th Annual Meeting of the North American*

- Chapter of the International Group for the Psychology of Mathematics Education.*
Roanoke, VA: PME-NA.
- Abrahamson, D., & Wilensky, U. (2006). ProbLab: An experimental unit on probability and statistics. *Manuscript in preparation.*
- Abrahamson, D., & Wilensky, U. (in press). Learning axes and bridging tools in a technology-based design for statistics. *International Journal of Computers for Mathematical Learning.*
- Benjamin, W. (1968). *Illuminations: Essays and reflections* (H. Zohn, Trans.). New York: Schocken.
- Cobb, P., Confrey, J., diSessa, A., Lehrer, R., & Schauble, L. (2003). Design experiments in educational research. *Educational Researcher*, 32(1), 9-13.
- diSessa, A. A. (2000). *Changing minds: Computers, learning and literacy*. Cambridge, MA: The MIT Press.
- Gelman, R., & Williams, E. (1998). Enabling constraints for cognitive development and learning: Domain specificity and epigenesis. In D. Kuhn & R. Siegler (Eds.), *Cognition, perception and language. Vol. 2* (5th ed., pp. 575-630). New York: John Wiley and Sons.
- Gigerenzer, G. (1994). Where do new ideas come from? In M. A. Boden (Ed.), *Dimensions of creativity* (pp. 53-74). Cambridge, MA: M.I.T. Press.
- Gigerenzer, G. (1998). Ecological intelligence: An adaptation for frequencies. In D. D. Cummins & C. Allen (Eds.), *The evolution of mind* (pp. 9-29). Oxford: Oxford University Press.
- Goldstone, R. I., & Wilensky, U. (2005). Promoting transfer through complex systems principles: Manuscript under review.
- Goldstone, R. I., & Wilensky, U. (under review). Promoting transfer through complex systems principles.
- Ifrah, G. (2000). *The universal history of numbers: From prehistory to the invention of the computer* (D. Bellos, E. F. Harding, S. Wood & I. Monk, Trans.). Toronto: John Wiley and Sons.
- Jay, M. (1988). Scopic regimes of modernity. In H. Foster (Ed.), *Vision and visibility* (pp. 3-23). New York: The New Press.
- Kuhn, T. (1962). *The structure of scientific revolutions*. Chicago: University of Chicago Press.
- Latour, B. (1990). Drawing things together. In M. Lynch & S. Woolgar (Eds.), *Representation in scientific practice* (pp. 19 – 68). Cambridge, MA: MIT Press.
- McLuhan, M. (1964). *Understanding media: The extensions of man*. New York: The New American Library.
- Melville, S., & Readings, B. (1995). *Vision and textuality*. Durham, NC: Duke University Press.
- Olson, D. R. (1994). *The world on paper*. Cambridge, UK: Cambridge University Press.
- Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas*. NY: Basic Books.
- Papert, S. (1991). Situating constructionism. In I. Harel & S. Papert (Eds.), *Constructionism* (pp. 1-12). Norwood, NJ: Ablex Publishing.
- Papert, S. (1996). A word for learning. In Y. Kafai & M. Resnick (Eds.), *Constructionism in Practice* (pp. 9-24). Mahwah, NJ: Lawrence Erlbaum Associates.
- Rotman, B. (2000). *Mathematics as sign: Writing, imagining, counting*. Stanford, CA: Stanford University Press.
- Saxe, G. B. (1981). Body parts as numerals: A developmental analysis of numeration among the Oksapmin in Papua New Guinea. *Child Development*, 52(1), 306-331.

- Saxe, G. B., & Esmonde, I. (in press). Studying cognition in flux: A historical treatment of fu in the shifting structure of Oksapmin mathematics. *Mind, Culture, & Activity*.
- Schmandt-Besserat, D. (1992). *Before writing Vol. 1: From counting to cuneiform*. Austin: University of Texas Press.
- Thagard, P. (1992). *Conceptual revolutions*. Princeton, NJ: Princeton University Press.
- The_Animated_Software_Company. (2006). Central limit theorem. In Internet glossary of statistical terms. from <http://www.animatedsoftware.com/statglos/sgcltheo.htm>
- Uttal, D. H., & DeLoache, J. s. (1997). Manipulatives as symbols: A new perspective on the use of concrete objects to teach mathematics. *Journal of Applied Developmental Psychology*, 18, 37-54.
- von Glasersfeld, E. (1987). Learning as a constructive activity. In C. Janvier (Ed.), *Problems of representation in the teaching and learning of mathematics* (pp. 3-18). Hillsdale, NJ: Lawrence Erlbaum.
- Weisstein, E. W. (2006a). Law of Large Numbers. In MathWorld--A Wolfram Web Resource. from <http://mathworld.wolfram.com/LawofLargeNumbers.html>
- Weisstein, E. W. (2006b). Probability. In MathWorld--A Wolfram Web Resource. from <http://mathworld.wolfram.com/Probability.html>
- Weisstein, E. W. (2006c). Statistics. In MathWorld--A Wolfram Web Resource. from <http://mathworld.wolfram.com/Statistics.html>
- Wilensky, U. (1993). *Connected mathematics—Building concrete relationships with mathematical knowledge* Unpublished manuscript, Cambridge, MA: M.I.T.
- Wilensky, U. (1995). Paradox, programming and learning probability. *Journal of Mathematical Behavior*, 14(2), 231-280.
- Wilensky, U. (1997). What is normal anyway?: Therapy for epistemological anxiety. *Educational Studies in Mathematics*, 33(2), 171-202.
- Wilensky, U. (1999). *NetLogo*. Northwestern University, Evanston, IL: The Center for Connected Learning and Computer-Based Modeling <http://ccl.northwestern.edu/netlogo/>.
- Wilensky, U. (2001). Modeling nature's emergent phenomena with multi-agent modeling languages. In *Proceedings of EuroLogo 2001*. Linz, Austria.
- Wilensky, U. (2006). *Complex systems and reformulating scientific disciplines: Implications for learning, analysis of social systems, and educational policy*. Paper presented at the annual meeting of the American Educational Research Association, San Francisco.
- Wilensky, U., & Papert, S. (2006). Restructurations: Reformulations of knowledge disciplines through new representational forms (Manuscript in preparation).
- Wilensky, U., Papert, S., Sherin, B., diSessa, A. A., Kay, A., & Turkle, S. (2005). *Center for Learning and Computation-Based Knowledge -- CLiCK. Proposal to the National Science Foundation, Science of Learning Center*. Unpublished manuscript.
- Wilensky, U., & Stroup, W. (1999a). *HubNet*. Evanston, IL: The Center for Connected Learning and Computer-Based Modeling, Northwestern University.
- Wilensky, U., & Stroup, W. (1999b). *Participatory simulations: Network-based design for systems learning in classrooms*. Paper presented at the Conference on Computer-Supported Collaborative Learning, Stanford University.
- Wolfram, S. (2002). *A new kind of science*. Champaign, IL: Wolfram Media, Inc.