Coming Soon to MoMath: Programmable Randomness

Jim Propp, UMass Lowell
MOVES 2025, August 11, 2025

Slides available at
http://faculty.uml.edu/jpropp/moves25a.pdf;
video available at
http://faculty.uml.edu/jpropp/moves25a.mp4.

Designing a new kind of quincunx for MoMath

Background: A quincunx (aka Galton board) is a device designed to illustrate the Central Limit Theorem: when many independent random variables contribute additively to some quantity, the distribution governing the sum tends toward a bell-shaped curve called a normal or Gaussian distribution.

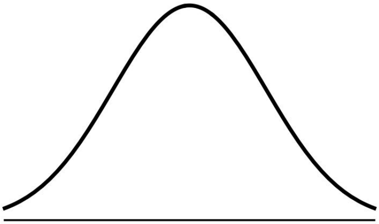

Designing a new kind of quincunx for MoMath

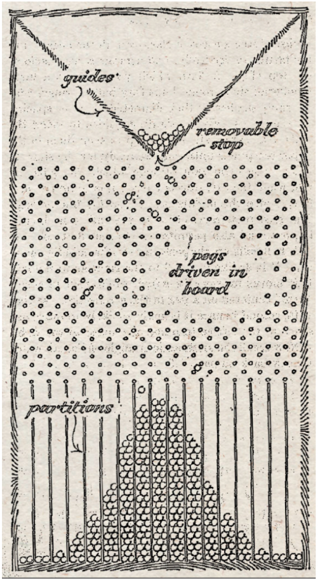

(From Warren Weaver's *Lady Luck: The Theory of Probability*, p. 261)

Designing a new kind of quincunx for MoMath

Assignment: Devise a quincunx in which the output distribution of the balls is not a centered binomial distribution (as is customary) or a shifted binomial distribution (as in the MoMath exhibit *Edge FX*) but an arbitrary user-specified distribution, by introducing individual biases at each of the junctions that route the balls from one level to the next (as was done in *Edge FX* for all the junctions in sync).

Exhibit name: *Draw Your Own Conclusions*

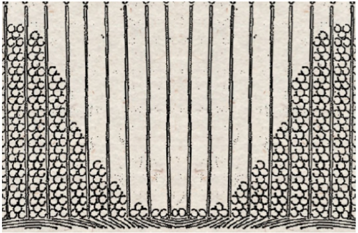

A binomial distribution (Galton)

This quincunx creates a binomial distribution on 3 bins:

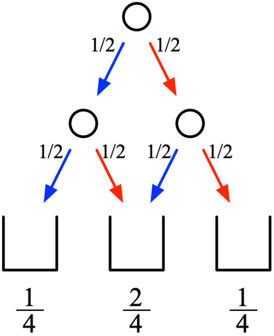

The number on an arrow leaving a junction signifies the proportion of balls, among those that arrive at the junction, that go in that direction.

A shifted binomial distribution (*Edge FX*)

This quincunx creates a shifted binomial distribution:

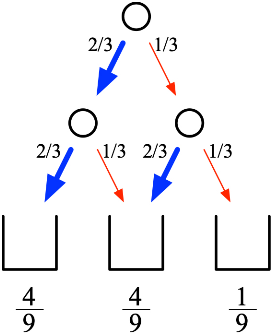

The probability that a ball follows a particular path is the product of the numbers along that path.

A non-binomial distribution (*D.Y.O.C.*)

This quincunx creates a uniform distribution on 3 bins:

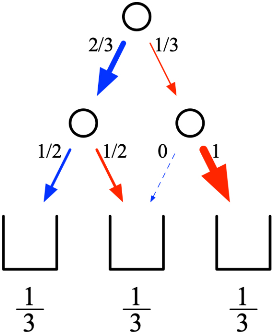

The probability that a ball arrives at a bin is the sum of the probabilities associated with all the paths to that bin.

Issues

Real world issues:

- Balls have momentum.
- Biases are hard to calibrate.
- Complicated parts break more quickly.

Mathematical issue:

- The problem is under-determined.

Too many variables, too few constraints

If we have $1 + 2 + 3 + \cdots + n$ junctions (with $n + 1$ bins at the bottom) then we have $\frac{n(n+1)}{2}$ biases to play with but only n degrees of freedom (we lose one degree of freedom because the probabilities associated with the bins must add up to 1).

Concrete example: the uniform distribution on 3 bins.
Our earlier solution has a mirror-twin.

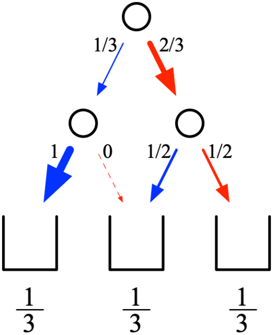

Introduce real-world desiderata

We'd like to minimize wear on the junctions to the extent possible.

It's too much to ask that all the junctions get used equally often; for one thing, every ball passes through the top junction, and that can't be true for all the other junctions!

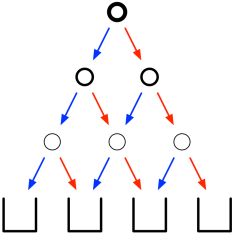

Introduce real-world desiderata

But we can require that, if you look at all the balls that end up in a particular bin, and see which path they took from the top of the quincunx to the bottom,

(*) ALL THE PATHS ARE TAKEN EQUALLY OFTEN.

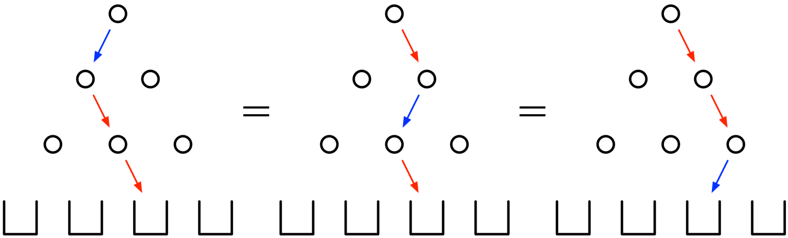

A solution for the uniform distribution on 3 bins

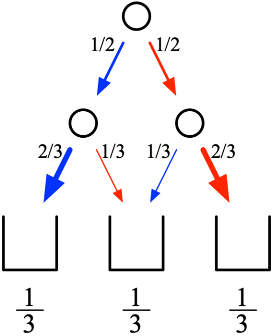

Here HALF the balls that wind up in the middle bin go left-then-right, and HALF the balls that wind up in the middle bin go right-then-left.

Property (*) is satisfied.

The general solution

Theorem: For each assignment of probabilities $p_0, p_1, \ldots, p_n$ to the $n + 1$ bins, there's a way to assign biases to the junctions so that the proportion of the balls that arrive (on average) at each of the $n + 1$ bins is equal to the probability assigned to that bin, AND so that property (*) is satisfied.

Moreover, outside of trivial cases where one or more of the bins is assigned probability zero, this assignment of biases is UNIQUE.

Furthermore, the biases are easy to compute through a two-part process of "image compression" and "image decompression".

To reduce resolution of a digital image ...

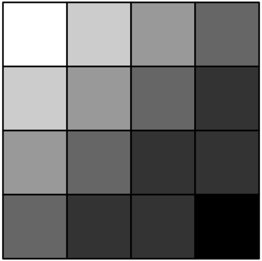

... introduce multi-pixel windows ...

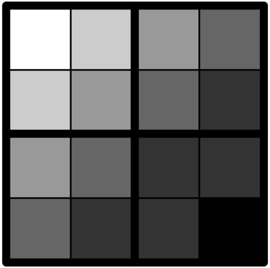

… and average the intensity within each windows

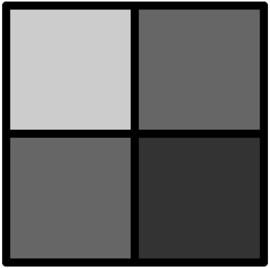

“Awkward” windows work too

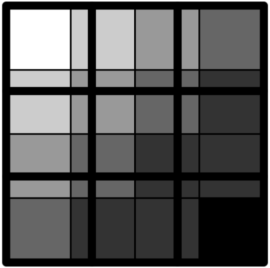

"Awkward" windows work too

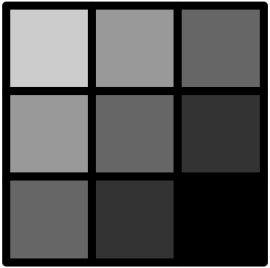

“Awkward windows work too

In my 2014 Gathering for Gardner (G4G11) talk, “The Programmable Galton Board: A Shameless Shill” I showed the compression process in reverse, starting from a single pixel and successively de-compressing it to a 2-by-2, 3-by-3, . . . image until I obtained a 256-by-256 picture of Tom Rodgers, the original sponsor of Gathering for Gardner.

https://www.youtube.com/watch?v=LDr8c2NmDDA

But what does this have to do with a quincunx?

Go down a dimension

Let's look at the averaging operation in one dimension. Here a, b, c, and d are the luminosity levels in a linear array containing 4 pixels (luminosity equals intensity times area). We compress down to 3 pixels with respective luminosities $(3a + 1b)/3$, $(2a + 2b)/3$, and $(1a + 3b)/3$.

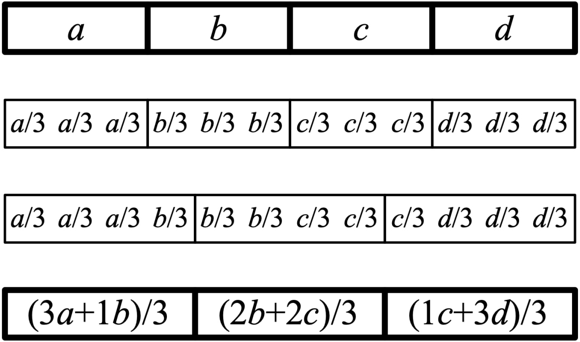

Go down a dimension

We can progressively compress from n pixels to $n-1$ pixels to $n-2$ pixels etc. down to just 1.

Why would we do this?

Because it can be shown that the resulting luminosity levels are exactly how we need probabilities to behave in our programmable quincunx!

Key insight: A quincunx isn't just giving a single distribution at the bottom; each level is governed by a distribution in which the "bins" are the pins!

Example: the classical quincunx

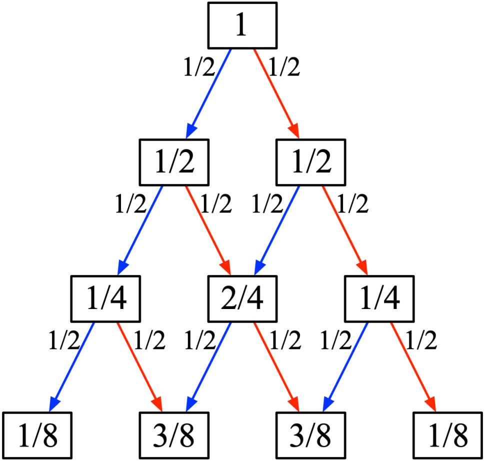

Example: the classical quincunx

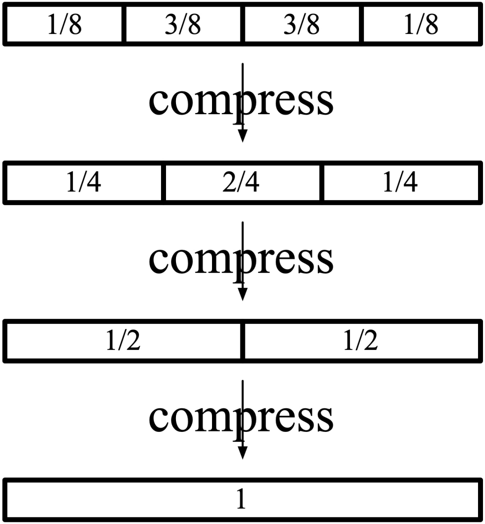

Example: a non-classical quincunx

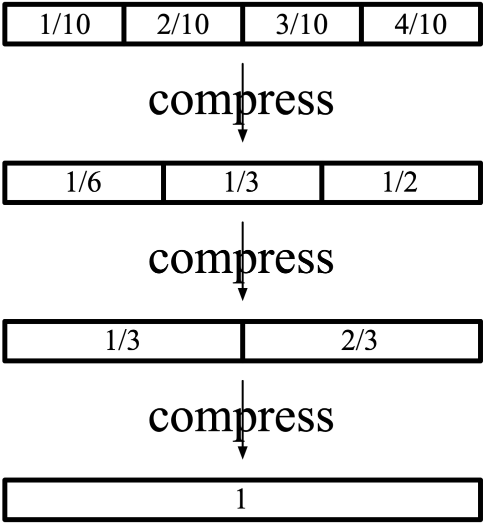

Example: a non-classical quincunx

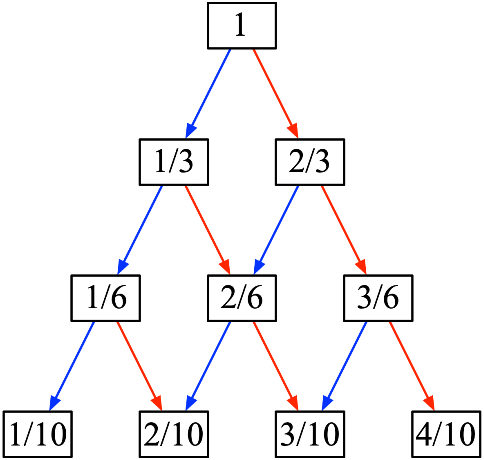

Example: a non-classical quincunx

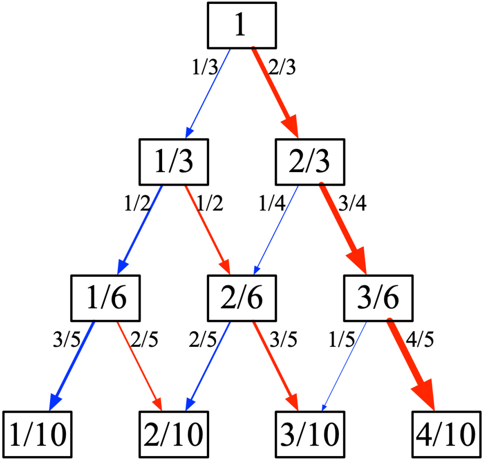

It lives! (virtually)

No one has come up with an economically viable way to build *Draw Your Own Conclusions* with pins, bins, and balls, but an electronic mock-up will debut at MoMath on Fifth in 2026.

Here's a screenshot of a preliminary version, shown with special permission:

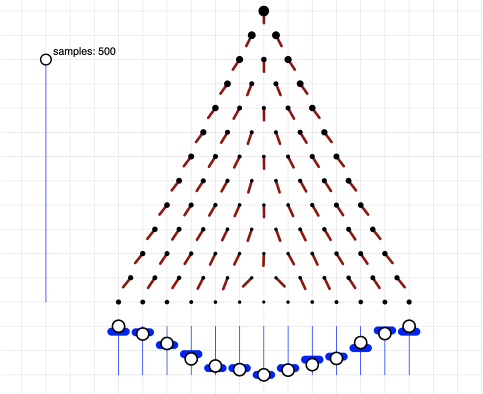

Keeping it real

What would it take to bring a real (i.e., physical) programmable quincunx to MoMath?

Many of the virtues of the scheme remain available even if one drastically limits the forms of bias at the junctions, allowing only THREE kinds of junctions: unbiased random junctions, nonrandom junctions that send all balls leftward, and nonrandom junctions that send all balls rightward.

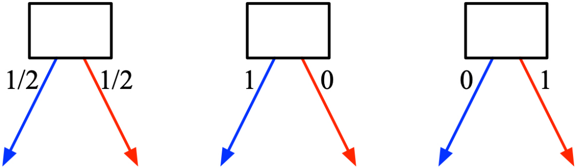

Keeping it real

This reduction in the complexity of the design would make the exhibit even cheaper to build and maintain.

It would introduce discrepancy in the distribution on the bins, but not much larger than the discrepancy due to randomness.

I'm hoping that the success of *Draw Your Own Conclusions* will inspire someone to fund a physical implementation.

In the meantime

I'd also be interested in variants that might be easier to build that don't use randomness.

For instance, instead of balls use water molecules. Junctions can easily divert some fraction of the water to the left or to the right. The resulting distribution in the buckets at the bottom exactly conform to the user-specified distribution.

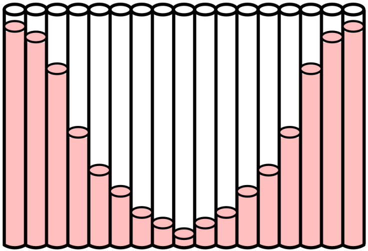

In the meantime

Alternatively, one could replace the junctions by deterministic "rotor routers" (generalizations of the sort of flipflops found in the old Think a Dot toy) that take turns sending balls leftward or rightward according to a fixed rotation schedule. This would deliver the same distribution on the balls at the bottom as a random quincunx but with less discrepancy.

Thanks for listening!

You can write to me at jamespropp@gmail.com.

If you want to financially support this exhibit, email Cindy Lawrence (executivedirector@momath.org).

Thank you!

Slides available at
http://faculty.uml.edu/jpropp/moves25a.pdf.