A Random Ladder Game: Permutations, Eigenvalues, and Convergence of Markov Chains

Lester H. Lange and James W. Miller

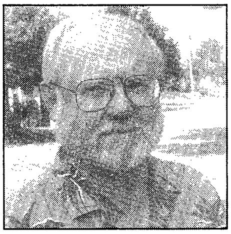

Les Lange is Emeritus Professor of Mathematics and Emeritus Dean of Science now serving San Jose State University as a volunteer at the Moss Landing Marine Laboratories on Monterey Bay. He learned mathematics at Berkeley, Valparaiso, Stanford, and Notre Dame. A complex variable Ph.D. thesis under Wladimir Seidel followed an M.S. with George Pólya. A Fellow of the California Academy of Sciences, he has written a linear algebra book; a geometry paper with Don Chakerian which received the MAA Lester R. Ford Sr. Award; and, with G. L. Alexanderson, the extensive Pólya obituary for the London Mathematical Society.

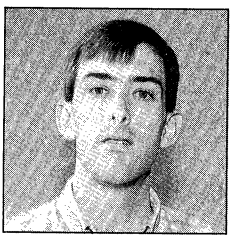

James W. Miller received a Bachelor of Music Education Degree from Baylor University in 1988. After teaching band for a short time, he determined that his interests were more in mathematics than music. He returned to Baylor and received an M.S. degree in Applied Mathematics in 1991. Most of the research for this paper was done during this time at Baylor. He is currently working on a Ph.D. degree at Southern Methodist University in the Department of Statistical Science. Primary research interests include matrix theory and stochastic processes.

1. A random ladder game: Another scheme that lets the gods decide. Last year a friend and his wife Reiko told how they'd each made a list of five things they'd like for Christmas, and then used an old scheme, which Reiko had learned as a little girl in Japan, to make a secret choice of which one item the recipient would receive as a gift from their spouse. It goes like this: The wife draws five vertical sticks, labelling them as "paths" A–E at the top. The numbers 1–5 (each secretly corresponding to one of the husband's gift requests) are written at the bottom of the respective paths. Then, in whatever pattern suits their fancy, they put into the figure a bunch of "rungs" between adjacent paths, with no rung sharing an endpoint with another rung, all of this as in Figure 1. The game proceeds according to this rule: The husband chooses a column heading. The wife moves down the path from the chosen point at the top; every time she encounters a rung she must cross it, then continue downward. Thus, for example, if "C" is chosen, she would end up at "2," and the husband would get the gift secretly corresponding to "2" for Christmas. The wife's gift is chosen similarly.

For the ladder shown in Figure 1, we have tabulated in Figure 2 the gift the husband would receive as a function of the column heading he chooses. The elements of the set $\{A, B, C, D, E\}$ have been mapped by this scheme into the set $\{1,2,3,4,5\}$. As the defining table shows, the map is a one-to-one, onto map. Observing this, we tried five or six different rung-path schemes—and it always happened that each of the elements in $\{1,2,3,4,5\}$ did occur as an image! Always so? For $\{1,2,3,\ldots,n\}$?... Reiko: "Of course. That's why it's used!" It's not hard to come up with a proof (see Section 3, below) that this ladder process does indeed

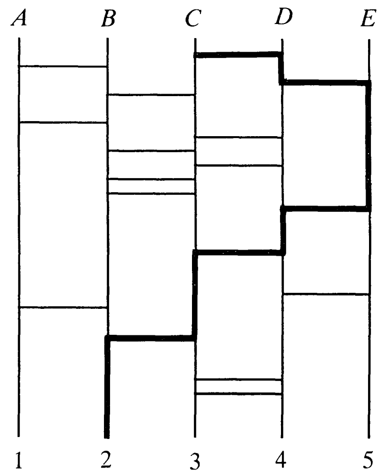

Figure 1
The path from C leads to 2.

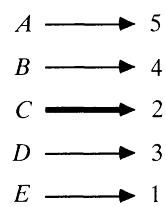

Figure 2

always yield such a one-to-one, onto map, a *permutation* of the set $\{1, 2, 3, \ldots, n\}$ when we think of the two involved sets as being the same.

There are other uses of the scheme. One of us went on to use it (a) to distribute four gifts randomly among four individuals at his office; and (b) to decide for his family who was to give whom a gift at Christmas.

2. More encounters. Sometime later, a colleague, Irene Cheung, who was raised in Hong Kong, said that she knew the scheme too. "We used it all the time in our office there to see who had to buy a treat for everybody at lunchtime." Participants would each write their names at the head of some path, put in a few rungs wherever they wished, and then let the next person do the same. Some player then wound up with the "losing" point of the day (selected independently). Irene said that in her experience it seemed to be in a player's interest to choose one of the outside paths—the ones headed A and E in our example—for the losing point was usually at the bottom of some interior path, and if you started along an outside path, it then seemed "you didn't usually wind up in the middle so far." We'll come back to this in Section 9.

The tale of yet another encounter begins in [**6**], where Martin Gardner begins by discussing a three-path scheme which decides for three computer programmers "who is to pay for the beer." What we've been calling "rungs" he calls "shuttles" —and his rules allow the shuttles to join those two outside paths. Gardner then draws figures for, enumerates, and names the six permutations that his ($n = 3$)

scheme generates, and proceeds to his principal interest for this column: introducing and discussing at some length the properties and structure of a group, in particular here, a permutation group. He points out that the group in his example has the very same structure as the group of transformations generated by the rotations and reflections of an equilateral triangle.

He then discusses "braiding three strands of a girl's hair," and also deals briefly with Emil Artin's theory of braids [**3**]. In this theory, the elements of the group are "weaving patterns" (infinite in number). The theory is involved in a game invented by Denmark's Piet Hein involving "tangled" braids, and in some theoretical physics questions dealt with by P. A. M. Dirac [**10**].

3. Some theorems. To prove that our ladder process produces a one-to-one map of a set $X = \{1, 2, 3, \ldots, n\}$ onto itself, let there be a finite number n of paths, and a finite number r of rungs, with no rung sharing an endpoint with another rung—all as in Figure 1, where $n = 5$. Now, assign one "marcher" to each of the n paths, and let them all start together and proceed abreast down their respective paths. Each time a rung is encountered, they pause while the two marchers who encountered the ends of that rung exchange paths. Then all start up again. There is always exactly one marcher per path, so each of the path-ends is reached by one and only one marcher. Thus, the map is indeed one-to-one and onto; a permutation of X.

Question: Does the process produce *all* possible permutations of X? If it does, then, among other things, we have the immediate corollary that *any permutation of a set X of n elements is a composition of transpositions* (of pairs of neighboring elements).

One way to prove that *all* permutations are generated by a ladder depends on this

Lemma. *By the insertion of certain rungs, any pair can by switched without disturbing the others.*

The key idea, which involves what we'll call a *transposition wedge*, is illustrated in Figure 3, where x and y get switched without disturbing the others. Such transposition wedges can be inserted to switch two elements as we please. This yields an obvious inductive proof of the result we want: if any particular permutation is specified, a wedge can move into the first spot whatever element you specify, and so on.

Here is a somewhat quicker proof of the fact that a ladder always yields a permutation of a set X of n elements. Let f be the function that a given ladder

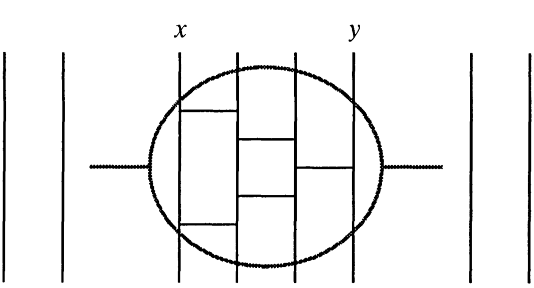

Figure 3
x and y get switched.

defines when it assigns to each of the elements in the domain (at the top) a definite element in the receiving set of n elements below. The rule of the game determines completely and unambiguously the way you come down the ladder to make that assignment. There is also only one way you can start at any given element below and climb up to the top, where you'll land on an element there. The ladder not only defines f, it also defines its inverse. This tells us that f must be one-to-one; that $f(X)$ consists of n distinct images; that all n points at the bottom of the ladder are in the game; and thus that f is indeed onto.

4. Markov chain models for a random ladder game. The Hong Kong office experiences reported in Section 2, lead to a doubt about the "fairness" of the ladder in producing a random permutation. To address the issue of fairness, we present two Markov chain models: *Model One* views the ladder as a random walk through the set of all permutations of a group of elements; *Model Two* considers the path of a single "marcher" (in the context of Section 3, above) as a random walk among the columns of the ladder. (Introductions to Markov chains abound; e.g., see [**2**]. For a more extensive treatment, see [**7**] and [**11**].)

5. Model One: A random walk about the set of all permutations of a set of *n* elements. Here, we define a Markov chain that changes states each time a rung is passed. The "state" of the process after the kth transition is the ordering of the n elements after passing the kth rung. To define this process formally, let $\{S(k),\ k = 0, 1, 2, \ldots, r\}$ be a stochastic process [**11**, p. 73] where $S(k)$ gives the state of the process after the kth rung has been passed. The state space for this process is the set of all possible permutations of n elements. Thus, there are $n!$ elements in the state space. Figure 4 illustrates how each rung of the ladder given

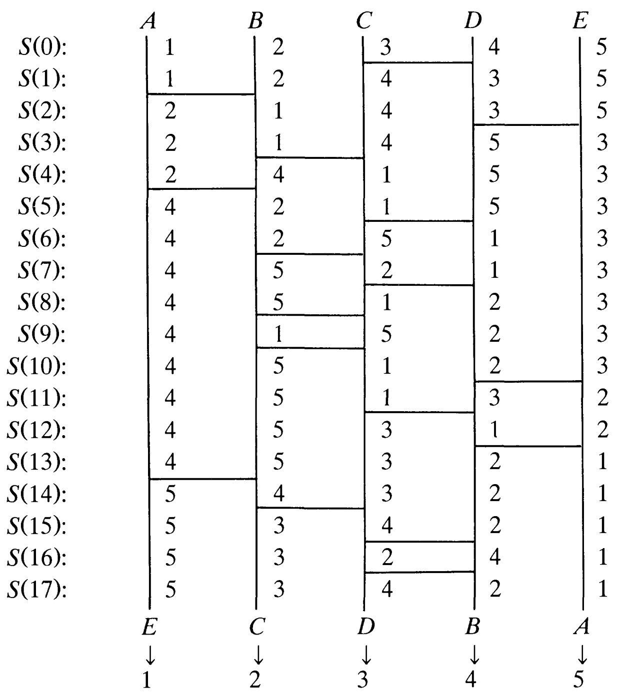

Figure 4

in Figure 1 changes the state of the process by transposing two neighboring elements.

We will assume the process begins with $S(0) = s_0$, the initial ordering $123\ldots n$. The remaining elements of the state space will be labelled arbitrarily as s_i, $i = 1, 2, \ldots, n! - 1$.

With the following assumptions, the process $\{S(k),\ k = 0, 1, 2, \ldots\}$ is a Markov chain:

(1) Assume that each time a rung is added to the ladder, it will be placed between the ith and $(i+1)$th columns with probability p_i, where $p_i > 0$ for $i = 1, 2, \ldots, n-1$ (so that the process is said to be "irreducible" [**11**, p. 142]—that is, it is possible to move from any state to any other state with appropriate rung placement) and $\sum_{i=1}^{n-1} p_i = 1$.

(2) Assume the rungs are placed independently.

Examples. When $n = 2$, there is only one possible position in which rungs may be placed. This process may be thought of as a walk back and forth between the two possible orderings "12" and "21." The *transition matrix* is

$$\mathbf{P} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix},$$

where the (i, j) entry is the probability of moving from state S_i to state S_j after the next rung. The transition diagram for the process is shown in Figure 5.

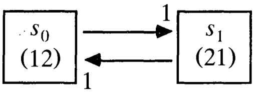

Figure 5

When $n = 3$, the transition diagram and transition matrix are shown in Figure 6. As the diagram shows, the $n = 3$ model may be thought of as a random walk about the nodes of a hexagon.

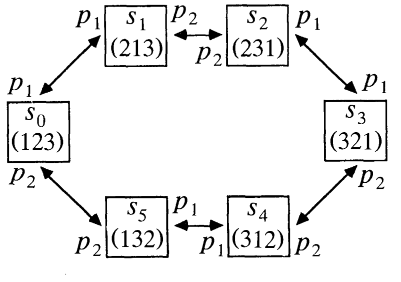

$$\mathbf{P} = \begin{bmatrix} 0 & p_1 & 0 & 0 & 0 & p_2 \\ p_1 & 0 & p_2 & 0 & 0 & 0 \\ 0 & p_2 & 0 & p_1 & 0 & 0 \\ 0 & 0 & p_1 & 0 & p_2 & 0 \\ 0 & 0 & 0 & p_2 & 0 & p_1 \\ p_2 & 0 & 0 & 0 & p_1 & 0 \end{bmatrix}.$$

Figure 6

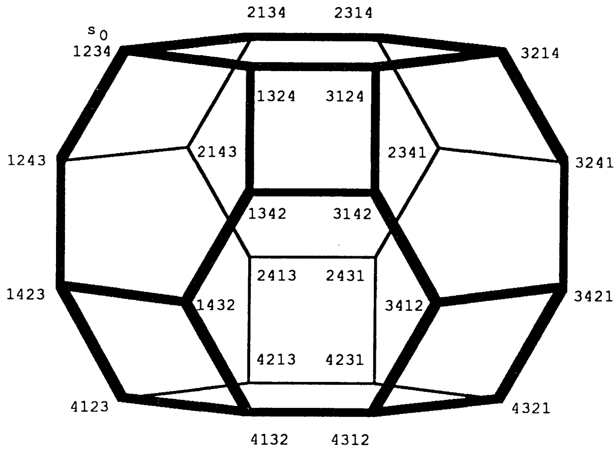

Figure 7

When $n = 4$, the model may be thought of as a random walk about the nodes of a truncated octahedron [**12**], an Archimedian solid consisting of 8 hexagonal faces and 4 square faces (see Figure 7).

Notice that our transition matrix is *stochastic*: *i.e.*, a square matrix with nonnegative entries such that the sum of the entries in each row is 1; its rows are probability vectors. In fact, since it is symmetric, it is *doubly stochastic*: its columns are also probability vectors [**4**, p. 48]. Crucial to the analysis of this process as a Markov chain is the fact that *for any choice of n, the transition matrix is symmetric, thus doubly stochastic*. We forego the easy proof that our transition matrices are symmetric; it depends on the fact that the only permutations that have a nonzero probability of being applied at any given step are those that transpose adjacent elements—and such permutations are their own inverses.

For a random ladder game, the transition matrix is easily seen to be doubly stochastic because it is symmetric. Feller makes the point in [**5**, p. 406] that *any* shuffling scheme yields a doubly stochastic transition matrix (even if it is not symmetric, as it is here). Since our ladder is a shuffling scheme, we don't really need to know that the transition matrix is symmetric in order to know that it is doubly stochastic. Still, the fact that it is symmetric is useful, for it tells us, thanks to an old theorem [**9**, p. 294], that its eigenvalues are all real. (See Section 8, below.)

6. Limiting probabilities. Let π_k be the row vector defined by

$$\pi_k^T = \begin{bmatrix} \Pr[S(k) = s_0] \\ \Pr[S(k) = s_1] \\ \vdots \\ \Pr[S(k) = s_{(n!-1)}] \end{bmatrix}.$$

The vector π_k gives for each state the probability that the process will be in that state after passing over k rungs. By assumption, $\pi_0 = [1 \quad 0 \quad 0 \quad \cdots \quad 0]$. This is the initial distribution vector, which places probability 1 on the event "the process begins in state s_0." For $n = 3$, using once more the designations in Figure 6, above, $\pi_1 = [0 \quad p_1 \quad 0 \quad 0 \quad 0 \quad p_2]$. This says that after one rung is passed, the process will be in state s_1 with probability p_1 (if the rung was between the first and second columns) or in state s_5 with probability p_2 (if the rung was between the second and third columns).

It can be shown that in general,

$$\pi_k = \pi_0 \mathbf{P}^k,$$

where $\mathbf{P}$ is the transition matrix for the process [**2**, p. 589; **11**, p. 139]. This provides a way of computing for each permutation the probability that it will be the result after k rungs have been passed. It is clear that each permutation is not equally likely after a small number of rungs. (For an easy example, the fact that $\pi_1 = [0 \quad p_1 \quad 0 \quad 0 \quad 0 \quad p_2]$ when $n = 3$ is proof that all outcomes are *not* equally likely when $n = 3$ and $k = 1$.)

Let us now consider $\pi_\infty = \lim_{k \to \infty} \pi_k$. Notice that

$$\pi_{k+1} = \pi_0 \mathbf{P}^{k+1} = \pi_0 \mathbf{P}^k \mathbf{P} = \pi_k \mathbf{P}.$$

Letting k go to infinity yields

$$\pi_\infty = \pi_\infty \mathbf{P}.$$

Thus, if this limit exists, it may be found by solving the linear system $\pi = \pi \mathbf{P}$ for $\pi = \pi_\infty$. This linear system does not have a unique solution, however, since any scalar multiple of π_∞ will also satisfy the system. Therefore, we need the added condition that the components of π_∞ sum to one, so that it is truly a probability vector. It can be shown [**11**, p. 151] that for irreducible Markov chains (such as our ladder), there is a unique solution to $\pi = \pi \mathbf{P}$ with components summing to one. If the limit π_∞ exists, it must be this unique solution.

For a random ladder game, the solution to $\pi = \pi \mathbf{P}$ with components summing to one is $\pi_\infty = (1/n!)\mathbf{u}$, where $\mathbf{u}$ is a row vector of ones. This fact may be proven as follows: Since the transition matrix is doubly stochastic, the elements of each column sum to one. This is equivalent to the matrix expression $\mathbf{uP} = \mathbf{u}$. The vector $\mathbf{u}$ is multiplied by $1/n!$ so that its elements sum to one. Since each component of π_∞ is the same, we conclude that *each ordering is equally likely as an outcome of the ladder in the limit as the number of rungs tends to infinity*.

We remind ourselves that the *existence* of π_∞ as the limit of π_k's has not been proven yet. As a matter of fact, $\lim_{k \to \infty} \pi_k$ *does not exist* because the process is periodic. As k tends to infinity, the sequence of π_k's approaches a *two-cycle*. That is, it approaches back-and-forth oscillation between two vectors we might call π_{even} and π_{odd}, where the ith component of π_{even} is $2/n!$ if s_i is an even permutation, 0 if s_i is an odd permutation; and the ith component of π_{odd} is $2/n!$ if s_i is an odd permutation, 0 if s_i is an even permutation. This fact is illustrated with the computed values of π_k given in Figure 8, which gives the values of π_k when there are three vertical columns and 0 through 10 horizontal rungs. This table also illustrates the fact that it is not necessary to assume that the rungs are placed uniformly on the ladder, that is, we are not assuming $p_i = 1/(n-1)$ for $i = 1, 2, \ldots, n-1$. We need only assume that these probabilities are positive and that the placements of the rungs are independent.

The statement "each ordering is equally likely as an outcome of the ladder in the limit as the number of rungs tends to infinity" still has validity for a random ladder game despite the periodicity of the process. We simply must interpret the statement in the following sense: The components of $\pi_\infty = (1/n!)\mathbf{u}$ give the

If $p_1 = p_2 = \frac{1}{2}$, then

$\pi_0 =$	[1	0	0	0	0	0]
$\pi_1 =$	[0	0.5	0	0	0	0.5]
$\pi_2 =$	[0.5	0	0.25	0	0.25	0]
$\pi_3 =$	[0	0.375	0	0.25	0	0.375]
$\pi_4 =$	[0.375	0	0.3125	0	0.3125	0]
$\pi_5 =$	[0	0.3438	0	0.3125	0	0.3438]
$\pi_6 =$	[0.3438	0	0.3281	0	0.3281	0]
$\pi_7 =$	[0	0.3359	0	0.3281	0	0.3359]
$\pi_8 =$	[0.3359	0	0.3320	0	0.3320	0]
$\pi_9 =$	[0	0.3340	0	0.3320	0	0.3340]
$\pi_{10} =$	[0.3340	0	0.3330	0	0.3330	0]

If $p_1 = \frac{1}{4}$, then $p_2 = \frac{3}{4}$, and

$\pi_0 =$	[1	0	0	0	0	0]
$\pi_1 =$	[0	0.25	0	0	0	0.75]
$\pi_2 =$	[0.625	0	0.1875	0	0.1875	0]
$\pi_3 =$	[0	0.2969	0	0.1875	0	0.5156]
$\pi_4 =$	[0.4609	0	0.2695	0	0.2695	0]
$\pi_5 =$	[0	0.3174	0	0.2695	0	0.4131]
$\pi_6 =$	[0.3892	0	0.3054	0	0.3054	0]
$\pi_7 =$	[0	0.3264	0	0.3054	0	0.3682]
$\pi_8 =$	[0.3578	0	0.3211	0	0.3211	0]
$\pi_9 =$	[0	0.3303	0	0.3211	0	0.3486]
$\pi_{10} =$	[0.3440	0	0.3280	0	0.3280	0]

Figure 8

long-run proportion of time that the process is in each state or the probability the outcome will be a particular state when the number of rungs is large and it is not known whether the number of rungs is even or odd.

The fact that each outcome becomes equally likely in the limit as k goes to infinity (see those bottom rows for π_9 and π_{10} in Figure 8) should not be too surprising. Imagine shuffling a deck of cards that was given to you with all cards initially in some particular order. Suppose you "shuffle" the deck by repeatedly selecting two adjacent cards at random and swapping their positions. Is it surprising that the deck becomes thoroughly shuffled after this process is applied for an indefinitely long time? This scheme is not the most efficient way to shuffle a deck, but after an indefinitely long time, the deck would get quite thoroughly shuffled.

Perhaps it *is* surprising that the outcomes become equally likely in the limit even when the rungs are not distributed uniformly throughout the ladder. Suppose for example that there is a very small probability that any given rung will appear between the first two columns. As long as there is a positive probability that a rung will appear there, eventually (after a sufficiently large number of rungs have been placed) a rung will appear there. As a matter of fact, eventually an arbitrarily large number of rungs will appear between the first two columns.

7. Model Two: The path of a single marcher as a random walk among the columns of the ladder. In Model One, above, we tracked all paths down the ladder simultaneously. The problem with that model is that the transition matrix is huge: $n! \times n!$. This makes the model very hard to work with computationally. For large values of n it is a difficult task even to write down the transition matrix.

Furthermore, finding its eigenvalues and evaluating its powers (of interest to us in Section 8, below) are even more difficult!

Model Two will trace the path of a *single* marcher moving down the ladder. The question considered here is: "Will a marcher beginning at the top of the ladder in column j be equally likely to end up in any of the n columns at the bottom of the ladder?" From the previous analysis, one could construct an argument to give the answer "virtually yes, if there is a sufficiently large number of rungs in the ladder." This second model will give the same result and provide a much simpler way to analyze the process.

Consider again a ladder with n vertical columns and r rungs. Let $\{T(k),\ k = 0, 1, 2, \ldots, r\}$ be a stochastic process where $T(k)$ gives the position (column number) of the marcher after the kth rung has been passed. The state space for this process is the set of all columns. Thus, there are only n elements in the state space. Again let p_i be the probability that any given rung lies between the ith and $(i+1)$th columns where $p_i > 0$ for $i = 1, 2, \ldots, n-1$ and $\Sigma_{i=1}^{n-1} p_i = 1$ and assume independence of the rung placements. Then this process is a Markov chain with just n states.

Each time the marcher passes a rung in the ladder, he will change positions to an adjacent column if the rung touches his column or he will remain in the same column if the rung does not touch his column. This model is essentially a random walk among the columns of the ladder. The transition diagram for this model and its transition matrix—an $n \times n$ tridiagonal matrix—are given in Figure 9.

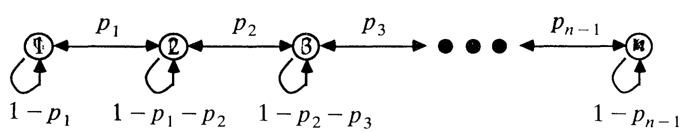

$$
\mathbf{P} = \begin{bmatrix}
1-p_1 & p_1 & 0 & 0 & & \cdots & 0 \\
p_1 & 1-p_1-p_2 & p_2 & 0 & & \cdots & 0 \\
0 & p_2 & 1-p_2-p_3 & p_3 & & \ddots & \vdots \\
0 & 0 & p_3 & & & \ddots & 0 \\
\vdots & \vdots & \ddots & \ddots & \ddots & & p_{n-1} \\
0 & 0 & \cdots & 0 & p_{n-1} & & 1-p_{n-1}
\end{bmatrix}
$$

Figure 9

This transition matrix is also symmetric, thus doubly stochastic for all n. It follows that the limiting probabilities are all $1/n$ just as expected. Regardless of where the marcher begins, he is equally likely (almost) to end up in any of the n columns at the bottom after a sufficiently large number of rungs has been passed. Unlike the first model, this process is not periodic for $n > 2$ since it is possible for the marcher not to change states when a rung is passed.

8. How many rungs is enough? Under each of the models given above, the transition matrix for the Markov chain is symmetric, thus doubly stochastic. This leads to the conclusion that each outcome becomes equally likely as the number of rungs grows to infinity. The natural question that arises at this point is: "How many rungs are needed in the ladder before each outcome is sufficiently equally likely?" Of course, "sufficiently equally likely" is subjective. What we are interested in is a way to determine how quickly k-step transition probabilities approach the limiting probabilities. This information lies in the eigenvalues of the transition matrix $\mathbf{P}$. This analysis can be applied to either of our two models.

To deal with the powers of the matrix $\mathbf{P}$, we need a few results about matrices: For any stochastic matrix $\mathbf{P}$, 1 is an eigenvalue, and *all* of its eigenvalues are contained in the closed unit circle in the complex plane [**4**, p. 49; **8**, p. 547]. The eigenvalues of symmetric matrices are always real [**9**, p. 294], so the eigenvalues of our transition matrices are real and lie in the interval $[-1,1]$. By *the small eigenvalues* of $\mathbf{P}$ we mean the eigenvalues of $\mathbf{P}$ whose absolute values are strictly less than 1. (This excludes 1 and -1.) By *the largest small eigenvalue* of $\mathbf{P}$ we mean the small eigenvalue having greatest absolute value. Another property of symmetric matrices that we shall use is the fact that there exists an orthogonal matrix $\mathbf{Q}$ such that $\mathbf{P}=\mathbf{QDQ}^T$, where the eigenvalues of $\mathbf{P}$ are the diagonal entries of the matrix $\mathbf{D}$ and all off-diagonal entries are zero [**9**, p. 297]. We can then more easily examine the powers of $\mathbf{P}$ by looking at the simpler powers of the diagonal matrix $\mathbf{D}$, since then $\mathbf{P}^k=\mathbf{QD}^k\mathbf{Q}^T$, which follows from the fact that, for orthogonal matrices, $\mathbf{Q}^T=\mathbf{Q}^{-1}$ [**9**, p. 230]. The powers of $\mathbf{D}$ are simply diagonal matrices containing powers of the eigenvalues of $\mathbf{P}$. The powers of the small eigenvalues will tend to zero as k goes to infinity. *How rapidly the sequence of π_k's approaches the limiting case depends on how rapidly the powers of the small eigenvalues approach zero.* Thus, the sequence will converge quickly when the largest small eigenvalue is close to zero. On the other hand, if there are small eigenvalues whose absolute values are close to one, the rate of convergence could be very slow.

Example. When $n=3$, the eigenvalues of the transition matrix in each case are not hard to find. In each model, the characteristic polynomial $\det\{\mathbf{P}-\lambda\mathbf{I}\}$ factors into the product of simple linear factors (corresponding to the roots 1 and perhaps -1) and quadratic factors. The quadratic formula can then be used to determine the remaining roots. These eigenvalues give insight into how quickly the outcomes become equally likely.

Under the first model, the transition matrix is given in Section 5, above. The eigenvalues of the transition matrix are $1, -1, \sqrt{p_1^2-p_1p_2+p_2^2}$ (multiplicity 2), and $-\sqrt{p_1^2-p_1p_2+p_2^2}$ (multiplicity 2). The absolute value of the largest small eigenvalue is $\sqrt{p_1^2-p_1p_2+p_2^2}$ which is minimized when $p_1=p_2=\frac{1}{2}$. The minimum value is then $\frac{1}{2}$. We may conclude that when $n=3$ and rung placement is uniform, the sequence $\{\pi_k,\ k=0,1,2,\dots\}$ converges to the two-cycle at roughly the same rate that the powers of $\frac{1}{2}$ converge to zero. We suspect (but have not shown) that for any n, the rate of convergence is maximized when rung placement is uniform. That is, the absolute value of the largest small eigenvalue is minimized when $p_i=1/(n-1)$ for $i=1,2,\dots,n-1$.

Under the second model, the transition matrix is

$$\mathbf{P}=\begin{bmatrix} p_2 & p_1 & 0 \\ p_1 & 0 & p_2 \\ 0 & p_2 & p_1 \end{bmatrix}$$

with eigenvalues 1, $\sqrt{p_1^2 - p_1 p_2 + p_2^2}$, and $-\sqrt{p_1^2 - p_1 p_2 + p_2^2}$. These are the same as the eigenvalues of the transition matrix when $n = 3$ under the first model except that in the first model, the multiplicities of the small eigenvalues were doubled and -1 was an eigenvalue. Could it be that something like this holds true for all n, so that the absolute values of the eigenvalues of the transition matrix under the first model are always the same as the absolute values of the eigenvalues of the transition matrix under the second model?

9. The best-case scenario. The conjecture was made in Section 8, above, that the rate of convergence to the limiting case is maximized when rung placement is uniform, i.e., $p_i = 1/(n-1)$ for all i. Assuming this conjecture is true, let us consider the rates of convergence under this best-case scenario for several different values of n. Figure 10 gives the value of the largest small eigenvalue e_n for $n = 3, 4, \dots, 10$ when the transition matrix is defined according to the second model. It can be seen that even in the best case, the largest small eigenvalue can be very close to one for large values of n. As a matter of fact, *all* of the eigenvalues approach the value 1 as n gets large. This occurs because the off-diagonal elements of the transition matrix $\mathbf{P}$ tend to zero, so the transition matrix approaches the identity matrix as n gets large. (Just see Figure 9 and remember that $p_i = 1/(n-1)$ here.) This means that even when rung placement is uniform, the convergence rate can be *very* slow. Figure 10 also gives the smallest integer r_n so that $e_n^{r_n} \leq 0.1$. This gives a rough indication of how many rungs are needed to "kill off" the largest small eigenvalue.

$n = 3$	$e_n = 0.5$	$r_n = 4$
$n = 4$	$e_n = 0.804739$	$r_n = 11$
$n = 5$	$e_n = 0.904508$	$r_n = 23$
$n = 6$	$e_n = 0.946410$	$r_n = 42$
$n = 7$	$e_n = 0.966990$	$r_n = 69$
$n = 8$	$e_n = 0.978251$	$r_n = 105$
$n = 9$	$e_n = 0.984923$	$r_n = 152$
$n = 10$	$e_n = 0.989124$	$r_n = 211$

Figure 10

Example. Consider the example in Section 2, above. Suppose six people use a ladder to determine who will have to buy everyone a treat at lunch. We look quickly at two cases: in the first case, we suppose that each of these six people places three rungs on the ladder at random, with rung placement uniform; in the second case, each participant places seven rungs. See Figure 11 which shows, for Model Two, the matrix $\mathbf{P}$ and both the 18-step and 42-step transition matrices. The ijth element of $\mathbf{P}^{18}$ gives the probability that a marcher beginning in column i will end up in column j after 18 rungs have been passed. It can be seen that if the marcher begins at one of the outside paths of the ladder, he will end up at the same position with probability 0.2866. Irene's comment (Section 2, above) that if you start on one of the outside paths, you don't end up in the middle so far is supported by this computation! For $n = 6$, 18 rungs are probably not sufficient to make the outcomes equally likely by most people's standards. On the other hand, if

$$\mathbf{P} = \begin{bmatrix} \frac{4}{5} & \frac{1}{5} & 0 & 0 & 0 & 0 \\ \frac{1}{5} & \frac{3}{5} & \frac{1}{5} & 0 & 0 & 0 \\ 0 & \frac{1}{5} & \frac{3}{5} & \frac{1}{5} & 0 & 0 \\ 0 & 0 & \frac{1}{5} & \frac{3}{5} & \frac{1}{5} & 0 \\ 0 & 0 & 0 & \frac{1}{5} & \frac{3}{5} & \frac{1}{5} \\ 0 & 0 & 0 & 0 & \frac{1}{5} & \frac{4}{5} \end{bmatrix}.$$

$$\mathbf{P}^{18} = \begin{bmatrix} 0.2866 & 0.2511 & 0.1931 & 0.1313 & 0.08221 & 0.05576 \\ 0.2511 & 0.2285 & 0.1893 & 0.1440 & 0.1048 & 0.08221 \\ 0.1931 & 0.1893 & 0.1795 & 0.1629 & 0.1440 & 0.1313 \\ 0.1313 & 0.1440 & 0.1629 & 0.1795 & 0.1893 & 0.1931 \\ 0.08221 & 0.1048 & 0.1440 & 0.1893 & 0.2285 & 0.2511 \\ 0.05576 & 0.08221 & 0.1313 & 0.1931 & 0.2511 & 0.2866 \end{bmatrix}$$

$$P^{42} = \begin{bmatrix} 0.1975 & 0.1892 & 0.1749 & 0.1584 & 0.1441 & 0.1359 \\ 0.1892 & 0.1832 & 0.1727 & 0.1606 & 0.1502 & 0.1441 \\ 0.1749 & 0.1727 & 0.1689 & 0.1645 & 0.1606 & 0.1584 \\ 0.1584 & 0.1606 & 0.1645 & 0.1689 & 0.1727 & 0.1749 \\ 0.1441 & 0.1502 & 0.1606 & 0.1727 & 0.1832 & 0.1892 \\ 0.1359 & 0.1441 & 0.1584 & 0.1749 & 0.1892 & 0.1975 \end{bmatrix}$$

Figure 11

each of the six participants were to place seven rungs on the ladder, the probabilities of the outcomes would be given by $\mathbf{P}^{42}$. Now the outcomes are pretty close to being equally likely. The problem, however, is this: after placing so many rungs on the ladder, it's a tedious task to trace the paths down!

The results prompted by Irene's comment led us to look at what happens if the ladder scheme is drawn on a cylinder, thus allowing easy transitions between the first and last columns. Suffice it to say that in this model, convergence rates were made faster, but as the number of columns increases, this effect diminishes.

10. Conclusion. A ladder shuffles the elements in a set, and as the number of rungs tends to infinity, the process does become fair: all the outcomes are equally likely. But to come close to that situation may require a number of rungs too large to be practical. Readers interested in other types of shuffling processes may wish to see [**5**, p. 406] and [**1**].

The ladder does provide a motivational setting for the pursuit by mathematics students of some of the important theorems of linear algebra. As for its usefulness in producing a random permutation or a random selection: Some people will surely go on yielding to the intriguing appeal of a random ladder. Others may prefer a more familiar method: drawing from a hat.

References

1. D. Aldous and P. Diaconis, Shuffling cards and stopping times, *American Mathematical Monthly* 93 (1986) 333–348.
2. H. Anton and C. Rorres, *Elementary Linear Algebra (Applications Version)*, 6th ed., Wiley, New York, 1991.

3. E. Artin, The theory of braids, *American Scientist* 38:1 (Jan 1950) 112–119.
4. A. Berman and R. J. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, Academic Press, New York, 1979.
5. W. Feller, *An Introduction to Probability Theory and Its Applications*, 3rd ed., Wiley, New York, 1968.
6. M. Gardner, Mathematical games, *Scientific American* (Dec. 1959) 166 ff. (This article also appears in his book, *Martin Gardner's New Mathematical Diversions from Scientific American*, New York, 1966.)
7. J. G. Kemeny and J. L. Snell, *Finite Markov Chains*, Van Nostrand, Princeton, 1960.
8. P. Lancaster, and M. Tismenetsky, *The Theory of Matrices*, 2nd ed., Academic Press, San Diego, 1985.
9. L. H. Lange, *Elementary Linear Algebra*, Wiley, New York, 1968.
10. M. H. A. Newman, On a string problem of Dirac, *The Journal of the London Mathematical Society*, Vol. 17, Part 3, No. 67 (July 1942) 173–177.
11. S. M. Ross, *Introduction to Probability Models*, 4th ed., Academic Press, San Diego, 1989.
12. G. L. Thompson, Exploratory Graphical Methods for Permutation Data, Southern Methodist University, Dallas, unpublished talk, April 1991.

The Half-Angle Formula for Cotangent

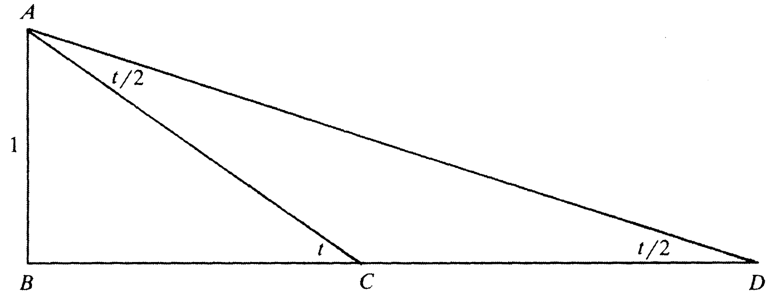

$ABCD$ is constructed so that $AB \perp BC$, $AC = CD$, and $AB = 1$. Then

$$m(\measuredangle CAD) = m(\measuredangle CDA) = 1/2m(\measuredangle ACB) = t/2,$$

$$\cot(t/2) = BC + CD = BC + AC$$

$$= \cot t + \sqrt{1 + \cot^2 t}.$$

While this construction is not new (see references), it is simple and efficient. This efficiency is not surprising since the compass construction of CD is a construction of the Carlyle circle through D.

References

1. H. L. Dorwart, Values of the trigonometric ratios of $\pi/8$ and $\pi/12$, in *Selected Papers on Precalculus*, ed. by Tom M. Apostol et al., MAA, Washington, DC, 1977, p. 195.
2. Roscoe Woods, The trigonometric functions of half or double an angle, in *Selected Papers on Precalculus*, ed. by Tom M. Apostol et al., MAA, Washington, DC, 1977, p. 189.
3. Duane W. DeTemple, Carlyle circles and the Lemoine simplicity of polygon constructions, *American Mathematical Monthly* 98 (1991) 97–108.

Submitted by Fen Chen, College Park, MD.