

ON THE RATE OF MIXING FOR p -SHUFFLES

STEVEN P. LALLEY

ABSTRACT. The p -shuffle is a natural generalization of the *dovetail* shuffle. It is defined as follows. First, the deck is cut into a top stack and a bottom stack so that the distribution of the size of the top stack is Binomial (N, p) , where N is the total number of cards in the deck. Then, conditional on the outcome of the cut, the two stacks are “riffled” in such a way that all possible riffles (interleavings) of these two stacks are equally likely. The main result of the paper is an asymptotic ($N \rightarrow \infty$) bound on the number of repetitions needed to “randomize” the deck.

1. INTRODUCTION

The p -shuffle (or *GSR- p* shuffle, for GILBERT, SHANNON, and REEDS) is a simple and natural generalization of the *dovetail* (or *GSR*) shuffle studied in [3] and [1]. For $0 < p < 1$ and a deck of size N , the p -shuffle is defined as follows: First, “cut” the deck into two stacks in such a way that the cardinality of the top stack has the Binomial (N, p) distribution. Next, “riffle” the two stacks by dropping cards one at a time from either the top stack or the bottom stack, according to the following rule: At any stage of the riffle, if there are A cards remaining in the top stack and B cards remaining in the bottom stack, then the probability that the next card dropped is from the top stack is $A/(A+B)$. The special case $p = 1/2$ is the dovetail shuffle of [3].

Mixing properties of the dovetail shuffle are elegantly set out by BAYER and DIACONIS in [3], following earlier work by ALDOUS [1]. There is a *cutoff phenomenon* for large decks: If the deck has N cards, then for large N the number of repetitions of the shuffle required for the distribution to approach to within total variation distance $1/e$ of the uniform distribution on all permutations is about $(3/2) \log_2 N$, and the transition to uniformity is rapid. (Recall that the *total variation distance* between two probability distributions μ, ν on a finite set \mathcal{X} is defined to be $(1/2) \sum_{x \in \mathcal{X}} |\mu(x) - \nu(x)|$.) In particular, if $d_N(n)$ is the total variation distance to the uniform distance after n repetitions, then for every $\varepsilon > 0$,

$$(1) \quad \lim_{N \rightarrow \infty} \inf_{n \leq ((3/2) - \varepsilon) \log_2 N} d_N(n) = 1, \text{ and}$$

$$(2) \quad \lim_{N \rightarrow \infty} \sup_{n \geq ((3/2) + \varepsilon) \log_2 N} d_N(n) = 0.$$

Bayer and Diaconis show that the transition to uniformity is even more rapid: see [3] for a precise statement.

It is natural to conjecture that a cutoff phenomenon holds in the p -shuffle, for any value of the parameter p . Moreover, it is tempting for the non-expert to suspect that routine

Date: January 29, 1999. Revised December 3, 1999.

1991 Mathematics Subject Classification. Primary: 60C05; Secondary: 60B15, 20B30.

Key words and phrases. Riffle shuffle, cutoff phenomenon.

Supported by NSF grant DMS-9626590.

modifications of the Aldous and Bayer/Diaconis analyses would establish the cutoff. Unfortunately, this is not the case. The dovetail shuffle has a special property not shared by any other p -shuffle: the probability distribution on the permutation group induced by one repetition of the dovetail shuffle is *uniform* on all “riffle permutations”. (A *riffle permutation* is defined to be a permutation with either one or two rising sequences; that is, a permutation which may result from one repetition of a p -shuffle.) This uniformity permits an exact formula for the distribution of the deck after n repetitions—see [3]. No such exact formula is available for the p -shuffle in general. Exact formulas for the eigenvalues of the transition probability matrices of the p -shuffles have recently been obtained – see [4] and [5]– but it does not seem that these can be used to determine the mixing rates.

The main result of this paper is an explicit bound on the rate of convergence for the p -shuffle. For each $p \in (0, 1)$, let $q = 1 - p$, and define constants C_p and C_p^* by

$$(3) \quad C_p := \frac{3 + \theta_p}{-4 \log(p^2 + q^2)} \text{ and}$$

$$(4) \quad C_p^* := \frac{2}{-\log(p^2 + q^2)},$$

where $\theta = \theta_p$ is the unique real number such that

$$(5) \quad p^\theta + q^\theta = (p^2 + q^2)^2.$$

Observe that $C_p < C_p^*$, and that $C_{1/2} = 3/(2 \log 2)$.

Theorem 1. *There exists an open neighborhood of $p = \frac{1}{2}$ such that for all p in this neighborhood, it takes at least $(C_p - \varepsilon) \log N$ and no more than $(C_p^* + \varepsilon) \log N$ repetitions of the p -shuffle to randomize a deck of size N . In particular, if $d_N(x)$ represents the total variation distance between the uniform distribution and the distribution of the state of the deck after $[x]$ repetitions of the shuffle, then for any $\varepsilon > 0$,*

$$(6) \quad \lim_{N \rightarrow \infty} d_N((1 - \varepsilon)C_p \log N) = 1 \text{ and}$$

$$(7) \quad \lim_{N \rightarrow \infty} d_N((1 + \varepsilon)C_p^* \log N) = 0.$$

The result (7) is due to JASON FULMAN [8]. (See also [7] for discussion of a related shuffling model.) A short proof of (7) is given in section 2.3 below—see Corollary 3. The relation (6) is considerably harder. Its proof is given in section 3. If the cutoff phenomenon does indeed occur for the p -shuffle, the cutoff must occur considerably before $C_p^* \log N$, at least when $p = 1/2$, since $C_{1/2}^*$ is strictly larger than $(3/2 \log 2)$. Thus, if the cutoff phenomenon occurs, C_p^* is *not* the right constant. We conjecture the right constant is C_p , at least for p sufficiently near $1/2$:

Conjecture 1. *For p sufficiently near $\frac{1}{2}$, there is a “cutoff phenomenon” at $C_p \log N$, that is (in light of (6)), for any $\varepsilon > 0$,*

$$(8) \quad \lim_{N \rightarrow \infty} d_N((1 + \varepsilon)C_p \log N) = 0.$$

A heuristic argument in favor of this conjecture is given in section 4. It seems likely that for *every* $p \in (0, 1)$ the cutoff phenomenon occurs at $C_p \log N$ for a suitable constant C_p , but our arguments are valid only for p near $\frac{1}{2}$. However, we will show (Proposition 1) that

for all $p > \frac{1}{2}$,

$$(9) \quad d_N \left((1 - \varepsilon) \frac{\log N}{\log p^{-1}} \right) = 1$$

for every $\varepsilon > 0$. Numerical analysis shows that for $p > .71$,

$$\frac{1}{\log p^{-1}} > \frac{3 + \theta_p}{-4 \log(p^2 + q^2)}.$$

Therefore, if for every $p \in [.5, 1)$ the cutoff phenomenon occurs at $C_p \log N$ for some constant C_p , and if our conjecture above is correct, then the function $p \mapsto C_p$ cannot be analytic in the parameter p .

Let $q = 1 - p$. There is an obvious “duality” between the p -shuffle and the q -shuffle: If one turns the deck upside down, then performs a p -shuffle, then turns the deck upside down again, the result is a q -shuffle. Thus, any result concerning the mixing rate of iterated p -shuffles holds also for q -shuffles. Henceforth, we shall assume that $1/2 \leq p < 1$.

2. PRELIMINARIES

2.1. The Associated Dynamical System. Imagine that the p -shuffle is performed infinitely many times, independently, on a deck of size N whose cards are initially labeled $1, 2, \dots, N$ from bottom to top. Each card i will have an *orbit*

$$x^i = x_1^i, x_2^i, \dots,$$

the sequence of 0s and 1s recording which stack (top=1, bottom=0) the card visits at each step. Note that the orbits of distinct cards are different, with probability 1, because on each repetition of the shuffle the conditional probability that cards i and j will be put in different stacks (given the history of the deck up to that repetition) is bounded below by a positive number.

Lemma 1. *The lexicographic order on the set of orbits $\{x^1, x^2, \dots, x^N\}$ coincides with the original order of the cards in the deck.*

Proof. This follows from the nature of the riffing process. If card i is below card j originally, then it will stay below for as long as the orbits of the two card coincide, because in each riffle cards are dropped one at a time from the bottoms of the two stacks. Thus, at the first time they are in *different* stacks, card i will be in the bottom stack (0) and card j will be in the top stack (1). Hence, $x^i \leq x^j$ in the lexicographic order. \square

Lemma 2. *The unordered collection of sequences x^1, x^2, \dots, x^N is a random sample of N independent Bernoulli- p processes. Equivalently, a version of the set $\{x^{(i)}\}_{1 \leq i \leq N}$ of orbits may be obtained by lex-ordering N independent Bernoulli- p sequences $\{x^i\}_{1 \leq i \leq N}$.*

Proof. See Diaconis [6], ch. 4D, for the case $p = \frac{1}{2}$; the general case is essentially the same. \square

The preceding lemmas imply that a version of the random walk on the permutation group S_N induced by independent repetitions of the p -shuffle may be (re-)constructed from the realizations of N independent Bernoulli- p sequences x^1, x^2, \dots, x^N . The position of the card i with orbit $x^{(i)}$ in the deck after n shuffles is determined as follows: apply the shift σ n times to each of the sequences x^j , and then determine the relative order of $\sigma^n x^{(i)}$ in the set $\{\sigma^n x^1, \dots, \sigma^n x^N\}$.

2.2. Example. In this example, the deck has $N = 5$ cards labelled a, b, c, d, e . The following table shows the arrangement of cards in the deck (from top to bottom) after $n = 0, 1, 2, 3, 4$ repetitions of the shuffle:

position	n=0	n=1	n=2	n=3	n=4
1	e	e	d	d	d
2	d	b	c	a	a
3	c	d	e	c	c
4	b	c	a	b	b
5	a	a	b	e	e

The next table the first 7 entries of each orbit:

e	0	0	0	1	0	1	1
d	0	1	0	0	0	0	0
c	0	1	0	1	0	1	0
b	1	0	1	1	0	0	1
a	1	1	1	0	0	0	1

Observe that the fifth entry of each orbit is 0, indicating that each card is in the “bottom” stack of the fourth riffle. Since the “top” stack is empty, the relative order of the cards is unchanged by the fourth riffle, that is, the fourth riffle permutation is the identity. Note that the same would be true if all the cards were in the “top” stack, that is, if the fifth entries of the orbits were all 1 instead of 0. This illustrates that the sequence of permutations does not uniquely determine the set of orbits.

2.3. The Associated Random Graph. The state of the deck after n repetitions of the p -shuffle may be described in terms of a random graph $\mathcal{G} = \mathcal{G}_n^N$ and an independent random permutation Π (uniformly distributed on the permutation group \mathcal{S}_N). This representation will facilitate comparison with the uniform distribution.

Let the orbits of the cards be $x^{(i)}$, $1 \leq i \leq N$, as described in section 2.1 above. Assume that the cards are labeled so that the original composition of the deck is $1, 2, \dots, N$, from bottom to top; thus, in the lexicographic order,

$$x^{(1)} \leq x^{(2)} \leq \dots \leq x^{(N)}.$$

Say that cards i, j have the same n -orbit if the orbits $x^{(i)}$ and $x^{(j)}$ coincide in their first n entries. Note that if i and j have the same n -orbit and $i < j$ then for every $i < i' < j$ the cards i and i' have the same n -orbit. Define $\mathcal{G} = \mathcal{G}_n^N$ to be the (random) graph with vertex set $[N] = \{1, 2, \dots, N\}$ and edge set \mathcal{E} consisting of those pairs $(i, i+1)$ such that cards i and $i+1$ have the same n -orbit. Observe that \mathcal{G} is completely determined by the set of n -orbits, i.e., by the $N \times n$ Bernoulli matrix

$$M = M^{(N,n)} = \begin{pmatrix} x_1^1 & x_2^1 & \cdots & x_n^1 \\ x_1^2 & x_2^2 & \cdots & x_n^2 \\ \cdots & \cdots & \cdots & \cdots \\ x_1^N & x_2^N & \cdots & x_n^N \end{pmatrix}$$

and consequently is independent of the (unordered) set of n -shifted orbits

$$\{\sigma^n x^{(i)} : i \in [N]\}.$$

Example: In the example described in section 2.2 above, the graph \mathcal{G}_4^5 has no edges, because the 4-orbits of the 5 cards are distinct. However, the graph \mathcal{G}_3^5 has an edge connecting vertices 2 and 3, because cards c and d have the same 3-orbit.

For a permutation π of $[N]$ define the \mathcal{G} -modification $\pi_{\mathcal{G}}$ of π to be the permutation obtained by ordering the assignments within each clique of \mathcal{G} . (A *clique* of \mathcal{G} is a maximal connected set of vertices.) Thus, for each singleton clique i (i.e., for each card i with a unique n -orbit),

$$\pi_{\mathcal{G}}(i) = \pi(i);$$

for each doubleton clique $i, i + 1$,

$$\begin{array}{llll} \pi_{\mathcal{G}}(i) = \pi(i) & \text{and} & \pi_{\mathcal{G}}(i + 1) = \pi(i + 1) & \text{if} & \pi(i) < \pi(i + 1) \text{ but} \\ \pi_{\mathcal{G}}(i) = \pi(i + 1) & \text{and} & \pi_{\mathcal{G}}(i + 1) = \pi(i) & \text{if} & \pi(i) > \pi(i + 1), \end{array}$$

and so on. Observe that, on the event that \mathcal{G} has no edges, $\pi_{\mathcal{G}} = \pi$.

Lemma 3. *Let $X_n = X_n^N$ be the permutation representing the state of the deck after n independent repetitions of the p -shuffle. (The i th entry $X_n(i)$ is the position of card i in the deck after n repetitions of the shuffle.) Let \mathcal{G} be the random graph associated to an $N \times n$ Bernoulli matrix M , and let Π be a random permutation independent of M with the uniform distribution on \mathcal{S}_N . Then the \mathcal{G} -modification of Π is a version of X_n .*

Proof. A version of the random walk $\{X_k\}_{k \geq 1}$ may be constructed from N independent Bernoulli- p processes x^1, x^2, \dots, x^N by the procedure described in section 2.1 above. Write

$$\begin{aligned} x^i &= z_1^i z_2^i \cdots z_n^i y_1^i y_2^i \cdots \\ z^i &= z_1^i z_2^i \cdots z_n^i \\ y^i &= y_1^i y_2^i \cdots \end{aligned}$$

The orbits of the cards $i = 1, 2, \dots, N$ are obtained by lex-ordering the N sequences x^i . This may be done in two steps: First, rearrange the rows by ordering the finite sequences z^i ; this may result in some ties. Second, rearrange the rows by ordering the infinite sequences y^i within any ties left by the first step. Since the rearrangement in the first step depends only on the values of the finite sequences z^i , the ordering of the infinite sequences y^i after the first stage is completely random. In particular, the assignment of row number i to the relative rank of the y -sequence in row i after the first step is a random permutation Π independent of the finite sequences z^i and uniformly distributed on the set \mathcal{S}_N . The permutation representing the deck after n p -shuffles is, clearly, obtained from Π by \mathcal{G} -modification, where \mathcal{G} is the random graph obtained from the Bernoulli matrix z . \square

Corollary 1. $d_N(n) \leq P\{\Pi \neq \Pi_{\mathcal{G}}\}$.

Proof. Recall that $d_N(n)$ is the total variation distance between the uniform distribution and the distribution of the deck after n repetitions of the p -shuffle. Since Π has the uniform distribution and $\Pi_{\mathcal{G}}$ has the same distribution as the random permutation resulting from n

repetitions of the p -shuffle, by the preceding lemma, it follows that

$$\begin{aligned}
d_N(n) &= \frac{1}{2} \sum_{\pi \in \mathcal{S}_N} |P\{\Pi_{\mathcal{G}} = \pi\} - P\{\Pi = \pi\}| \\
&= \frac{1}{2} \sum_{\pi \in \mathcal{S}_N} |P\{\Pi_{\mathcal{G}} = \pi \text{ and } \Pi \neq \Pi_{\mathcal{G}}\} - P\{\Pi = \pi \text{ and } \Pi \neq \Pi_{\mathcal{G}}\}| \\
&\leq \frac{1}{2} \sum_{\pi \in \mathcal{S}_N} P\{\Pi_{\mathcal{G}} = \pi \text{ and } \Pi \neq \Pi_{\mathcal{G}}\} + \frac{1}{2} \sum_{\pi \in \mathcal{S}_N} P\{\Pi = \pi \text{ and } \Pi \neq \Pi_{\mathcal{G}}\} \\
&= P\{\Pi \neq \Pi_{\mathcal{G}}\}.
\end{aligned}$$

□

Define

$$(10) \quad \tau = \tau_N := \min\{n : \text{no two rows of } M^{(N,n)} \text{ are the same}\}.$$

Since the rows of $M^{(N,\infty)}$ are independent sequences of i.i.d. Bernoulli- p random variables, τ_N is almost surely finite. Clearly, if $n \geq \tau_N$ then the random graph $\mathcal{G} = \mathcal{G}_n^N$ has no edges, and so $\Pi_{\mathcal{G}} = \Pi$. This implies

Corollary 2. $d_N(n) \leq P\{\tau_N > n\}$.

Observe that τ_N is a *strong stationary time* in the sense of [2]. It is likely that the bounds on $d_N(n)$ obtained via Corollary 2 are loose.

Corollary 3. $\lim_{N \rightarrow \infty} d_N((1 + \varepsilon)C_p^* \log N) = 0$.

Proof. The entries of the matrix $M^{(N,n)}$ are i.i.d. Bernoulli- p random variables; thus, the probability that any two given rows are identical is $(p^2 + q^2)^n$. It follows that the *expected* number of pairs of identical rows is

$$\binom{N}{2} (p^2 + q^2)^n < N^2 (p^2 + q^2)^n.$$

If $n \geq (1 + \varepsilon)C_p^* \log N$, then $(p^2 + q^2)^n \leq N^{-2(1 + \varepsilon)}$. Consequently, by the Chebyshev–Markov inequality, for $n = \lceil (1 + \varepsilon)C_p^* \log N \rceil$,

$$P\{\tau_N > n\} = P\{M^{(N,n)} \text{ has two identical rows}\} \leq N^{-2\varepsilon} \longrightarrow 0$$

as $N \rightarrow \infty$.

□

2.4. Entropy and Information. The construction of the preceding section shows that the distribution of the permutation X_n representing the state of the deck after n p -shuffles is completely determined by the distribution of the random graph \mathcal{G} . When $p \neq \frac{1}{2}$ this distribution may be rather complicated; for instance, when $p > \frac{1}{2}$ cliques are more likely in certain “hot spots” of the deck where n -orbits are more likely to contain more than the average number np of 1s. This is because when $p \neq \frac{1}{2}$ not all n -orbits are equally likely, and, in fact, the likelihoods of different n -orbits may be on entirely different orders of magnitude. It will ultimately be necessary to estimate the number of different n -orbits with likelihood on a given order of magnitude. These estimates will involve *entropy* and *information* numbers, about which we gather some basic information here.

For an n -orbit $x = x_1x_2 \dots x_n$ (a 0-1 sequence of length n) let $\bar{x} = \sum x_i/n$ be the relative frequency of 1s. For $\alpha \in (0, 1)$, $\varepsilon > 0$, and $k \leq n$ define

$$\begin{aligned} \Omega_n &= \{0, 1\}^n \quad \text{and} \\ \Omega_n^\varepsilon(\alpha) &= \{x \in \Omega_n : |\bar{x} - \alpha| < \varepsilon\}. \end{aligned}$$

Define the *Shannon entropy function* $H(\beta)$ and the *Kullback–Leibler information function* $I(p, \beta)$ by

$$\begin{aligned} H(\beta) &= -\beta \log \beta - (1 - \beta) \log(1 - \beta) && \text{for } 0 < \beta < 1, \\ I(p, \beta) &= -\beta \log p - (1 - \beta) \log q && \text{for } 0 < \beta, p < 1. \end{aligned}$$

The following well-known lemma is an easy consequence of Stirling's formula.

Lemma 4. *For each $\alpha \in (0, 1)$ and each $\varepsilon > 0$ small enough that $0 < \alpha - \varepsilon < \alpha + \varepsilon < 1$,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log |\Omega_n^\varepsilon(\alpha)| = \max_{|\beta - \alpha| \leq \varepsilon} H(\beta).$$

Fix $p \in (\frac{1}{2}, 1)$. For each sequence $x = x_1x_2 \dots x_n \in \Omega_n$, define the *likelihood*

$$\lambda(x) = \prod_{i=1}^n (p^{x_i} q^{1-x_i}).$$

Corollary 4. *For each $\alpha \in (0, 1)$ and $\varepsilon > 0$, and for every fixed integer $k \geq 1$,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \left(\sum_{x \in \Omega_n^\varepsilon(\alpha)} \lambda(x)^k \right) = \max_{|\beta - \alpha| \leq \varepsilon} (H(\beta) - kI(p, \beta))$$

Proof. This is a routine consequence of Lemma 4 and the continuity of the functions H and I in their arguments. \square

Fix $p \in (\frac{1}{2}, 1)$, and for $t \geq 0$ define

$$(11) \quad \psi(t) = \log(p^t + q^t),$$

$$(12) \quad p_t = p^t / (p^t + q^t),$$

$$(13) \quad q_t = q^t / (p^t + q^t) = 1 - p_t.$$

Since $p > \frac{1}{2}$, the function $t \mapsto p_t$ is strictly increasing in t , with $p_0 = \frac{1}{2}$ and $p_\infty = 1$. Moreover, $\psi(t)$ is strictly decreasing in t , with $\psi(0) = \log 2$ and $\psi(\infty) = -\infty$. Consequently, since $\psi(2) < 0$, there is a unique $\theta > 0$ such that $\psi(\theta) = 2\psi(2)$ —this is the value singled out in the statement of Theorem 1. Notice that for every $t > 0$,

$$(14) \quad \psi'(t) = p_t \log p + q_t \log q = -I(p, p_t)$$

where $I(\cdot, \cdot)$ is the Kullback–Leibler information function. Since p_t is increasing in t , $I(p, p_t)$ is decreasing in t , and therefore $\psi(t)$ is strictly convex in t . The functions H, I , and ψ are related by the identities

$$(15) \quad H(p_t) = tI(p, p_t) + \psi(t) = -t\psi'(t) + \psi(t),$$

which shows that $-H(p_t)$, considered as a function of $-I(p, p_t)$, is the Legendre transform of $\psi(t)$.

3. PROOF OF THEOREM (1)

We have already proved the easy inequality (7) in Corollary 3 above. It remains to prove the hard inequality (6). Fix $p > \frac{1}{2}$ and $C < C_p$, and let

$$(16) \quad n = \lfloor C \log N \rfloor$$

be the number of repetitions of the p -shuffle. As usual, N is the cardinality of the deck. In this section we will show that if $p - \frac{1}{2}$ is sufficiently small then

$$(17) \quad \lim_{N \rightarrow \infty} d_N(n) = 1.$$

Since the total variation distance $d_N(n)$ is, for fixed N , nonincreasing in n , we may assume in proving (17) that C is arbitrarily close to C_p . Recall that $C_{1/2} = 3/(2 \log 2)$, so for p near $\frac{1}{2}$, C_p is close to $3/(2 \log 2)$. Consequently, if $\beta > 0$ is small then for p near $\frac{1}{2}$ and C near C_p ,

$$(18) \quad p^n < N^{-1-\beta}.$$

3.1. Clumping. There are (at least) two obstructions to mixing for p -shuffles, *clumping* and *cold spots*. The term *clumping* refers to a tendency for large numbers of cards at the top or bottom of the deck to remain in their original order. This may occur in the p -shuffle for p near 0 or 1 because the deck is, with high probability, so unevenly divided that cards in one of the two stacks do not have a chance to be mixed with cards from the other. The term *cold spots* refers to the existence of predictable areas of the deck in which unusually large numbers of neighboring *pairs* $i, i+1$ tend to remain in their original order. For p near $\frac{1}{2}$, the existence of cold spots is the obstruction that persists for larger values of n , while for $p > .71$ clumping is the persistent obstruction. The following lemma shows that clumping will occur when $C < 1/\log p^{-1}$.

Proposition 1. *For any $p \in [\frac{1}{2}, 1)$, if $C < 1/\log p^{-1}$ then $\lim_{N \rightarrow \infty} d_N(n) = 1$.*

Proof. If $C < 1/\log p^{-1}$, then there exists $\varepsilon > 0$ such that $p^n \geq N^{-1+\varepsilon}$ for all N . But p^n is the probability that a randomly chosen card will have n -orbit $111 \cdots 1$. Since the N orbits are independent, it follows that, with probability approaching 1 as $N \rightarrow \infty$, the number of cards with n -orbit $111 \cdots 1$ is at least $N^{\varepsilon/2}$. On this event, the $\lfloor N^{\varepsilon/2} \rfloor$ topmost cards of the deck remain in their original order after n shuffles. But for a completely random (i.e., uniformly distributed) permutation, the length of the longest run of consecutive cards that remain in their original order is, with probability approaching 1 as $N \rightarrow \infty$, smaller than $(\log N)^2$. (To see this, observe that the probability that $2k$ consecutive cards remain in order is smaller than 2^{-k} , because, in a completely random permutation, the probability that two neighboring cards remain in order is $1/2$, and these events are independent for nonoverlapping pairs.) \square

3.2. Cold Spots. The existence or nonexistence of “cold spots” in the deck depends on the statistics of the random graph \mathcal{G} associated to the set of n -orbits (see section 2.3). Let \mathcal{E} be the edge set of \mathcal{G} ; define

$$\mathcal{K} = \{i \in [N] : (i, i+1) \in \mathcal{E}\}.$$

For a subset $H \subset [N]$ let ∂H be the (inner) *boundary* of H in $[N]$, i.e., the subset of H consisting of all points at distance 1 from $[N] - H$.

Proposition 2. *Suppose that there exist nonrandom subsets $H = H_N$ of $[N]$ such that for some $\varepsilon > 0$ all three of the following conditions are met as $N \rightarrow \infty$:*

$$(19) \quad |H| \longrightarrow \infty;$$

$$(20) \quad |\partial H| = O(|H|^{\frac{1}{2}}); \text{ and}$$

$$(21) \quad P\{|\mathcal{K} \cap H| \geq |H|^{\frac{1}{2}+\varepsilon}\} \rightarrow 1.$$

Then

$$\lim_{N \rightarrow \infty} d_N(n) = 1.$$

Note: The sets $H = H_N$ are the ‘‘cold spots’’. Aldous [1] states the special case $p = \frac{1}{2}$ and $H = [N]$, but his proof appears to have serious gaps.

To prove Proposition 2, we will show that under the hypotheses (19), (20), and (21), a recognizably large number of pairs of neighboring cards in H_N remain in their original order after n repetitions of the shuffle. For any permutation π and an index $i \in [N - 1]$ define $\zeta_i = \zeta_i(\pi)$ by

$$(22) \quad \begin{aligned} \zeta_i &= 1 && \text{if } \pi(i) < \pi(i + 1), \\ \zeta_i &= 0 && \text{if } \pi(i) > \pi(i + 1) \end{aligned}$$

The number of *alignments* in H is defined to be $\sum_{i \in H} \zeta_i$. Let Π be a completely random permutation, and let $\Pi_{\mathcal{G}}$ be the \mathcal{G} -modification of Π . Observe that the random variables $\zeta_i(\Pi)$ are Bernoulli- $\frac{1}{2}$, so the expected number of alignments in H is $|H|/2$.

Lemma 5. *For any interval $J = \{a, a + 1, \dots, b\} \subset [N - 1]$,*

$$E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}}) \mid \mathcal{G} \right) \geq \frac{1}{2}|J| + \frac{1}{6}|\mathcal{K} \cap J| - 3.$$

Proof. For a fixed graph $\mathcal{G} = ([N], \mathcal{E})$ whose vertex set is $[N]$ and whose edge set \mathcal{E} is contained in $\{(i, i + 1) : 1 \leq i \leq N - 1\}$, define a sequence of subgraphs $\mathcal{G}_k = ([N], \mathcal{E}_k)$ by taking for the edge set \mathcal{E}_k the set of all edges $(i, i + 1) \in \mathcal{E}$ such that $i \leq k$. The edge sets of these graphs are nested, with $\mathcal{G}_0 = ([N], \emptyset)$ and $\mathcal{G}_N = \mathcal{G}$, and \mathcal{G}_{k+1} is obtained from \mathcal{G}_k by adjoining *at most* the single edge $(k + 1, k + 2)$ to the edge set of \mathcal{G}_k .

A completely random permutation Π may be constructed by attaching independent, identically distributed uniform-(0,1) random variables U_i to the indices $i = 1, 2, \dots, N$ and defining $\Pi(i)$ to be the relative rank of U_i in the set $\{U_j\}_{1 \leq j \leq N}$. The \mathcal{G}_k -modification of Π is obtained by first permuting the r.v.s U_i within \mathcal{G}_k -cliques so that the U 's within cliques are ordered, then setting $\Pi_{\mathcal{G}_k}(i)$ to be the relative rank of the U attached to the index i . Consider first the graph \mathcal{G}_0 . Since its edge set is empty, the \mathcal{G}_0 -modification of Π is identical to Π ; consequently,

$$E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}_0}) \mid \mathcal{G}_0 \right) = \frac{1}{2}|J|.$$

Now consider the change in $\zeta_i(\Pi_{\mathcal{G}_k})$ when k is incremented to $k + 1$. For those k such that $k + 2 < a$ (a is the leftmost index in J), none of the values $\zeta_i, i \in J$, is affected, since

adding (at most) the single edge $(k+1, k+2)$ does not change the U attached to any $i \in J$. Hence, if $k+2 < a$, then

$$E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}_{k+1}}) \middle| \mathcal{G}_{k+1} \right) = E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}_k}) \middle| \mathcal{G}_k \right).$$

Next, consider $k \geq b$ (b is the rightmost index in J). Adding the edge $(k+1, k+2)$ changes only the rightmost clique; consequently, the only value $\zeta_i, i \in J$, that may be affected is at the i immediately to the left of the rightmost clique in J . Hence,

$$\left| E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}_N}) \middle| \mathcal{G}_N \right) - E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}_{b-1}}) \middle| \mathcal{G}_{b-1} \right) \right| \leq 1.$$

Similarly, if $k = a-2$ or $k = a-1$, the only values of $\zeta_i, i \in J$, that may be affected are ζ_a and ζ_{a+1} , and so

$$\left| E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}_{a-1}}) \middle| \mathcal{G}_{a-1} \right) - E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}_0}) \middle| \mathcal{G}_0 \right) \right| \leq 2.$$

Finally, consider the changes in the values ζ_i when k is incremented to $k+1$ for $a \leq k < b$. If the edge $(k+1, k+2)$ is *not* added, then obviously nothing is changed. If $(k+1, k+2)$ *is* added, then at most three values ζ_i are affected: ζ_{k+1}, ζ_{k+2} , and ζ_i for the index i immediately to the left of the clique containing the vertex $k+1$. Let c be the size of the \mathcal{G}_k -clique containing $k+1$, and let c' be the size of the \mathcal{G}_k -clique containing i . Conditional on \mathcal{G}_k and \mathcal{G}_{k+1} , respectively,

$$\begin{aligned} \zeta_i(\Pi_{\mathcal{G}_k}) &\sim \text{Bernoulli-}(1/\binom{c+c'}{c}), & \zeta_i(\Pi_{\mathcal{G}_{k+1}}) &\sim \text{Bernoulli-}(1/\binom{c+c'+1}{c+1}) \\ \zeta_{k+1}(\Pi_{\mathcal{G}_k}) &\sim \text{Bernoulli-}(1/(c+1)), & \zeta_{k+1}(\Pi_{\mathcal{G}_{k+1}}) &\sim \text{Bernoulli-}1 \\ \zeta_{k+2}(\Pi_{\mathcal{G}_k}) &\sim \text{Bernoulli-}(1/2), & \zeta_{k+2}(\Pi_{\mathcal{G}_{k+1}}) &\sim \text{Bernoulli-}(1/(c+2)) \end{aligned}$$

Thus, the net change in (conditional) expectation is

$$\begin{aligned} (23) \quad & E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}_{k+1}}) \middle| \mathcal{G}_{k+1} \right) - E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}_k}) \middle| \mathcal{G}_k \right) \\ &= 1 - \frac{1}{2} + \frac{1}{c+2} - \frac{1}{c+1} + \frac{1}{\binom{c+c'+1}{c+1}} - \frac{1}{\binom{c+c'}{c}} \\ &\geq \frac{1}{2} - \frac{1}{6} - \frac{1}{6} = \frac{1}{6}. \end{aligned}$$

Consequently, the conditional expectation is incremented by at least $1/6$ for every edge $(k+1, k+2)$ of \mathcal{G} in J , so

$$E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}_{b-1}}) \middle| \mathcal{G}_{b-1} \right) - E \left(\sum_{i \in J} \zeta_i(\Pi_{\mathcal{G}_{a-1}}) \middle| \mathcal{G}_{a-1} \right) \geq \frac{1}{6} |J \cap \mathcal{K}|.$$

□

Lemma 6. *For any set $H \subset [N - 1]$,*

$$\text{var} \left(\sum_{i \in H} \zeta_i(\Pi_{\mathcal{G}}) \middle| \mathcal{G} \right) \leq 3|H|$$

Proof. Let $\Gamma_1, \Gamma_2, \dots, \Gamma_m$ be the cliques of \mathcal{G} , listed in order from left to right. Each index i is in a unique clique Γ_j , so

$$\sum_{i \in H} \zeta_i = \sum_{j=1}^m Z_j$$

where

$$Z_j = \sum_{i \in H \cap \Gamma_j} \zeta_i.$$

Within any clique, all ζ_i must be 1 except for the rightmost index i , so $Z_j - |H \cap \Gamma_j| + 1$ is a Bernoulli random variable. Furthermore, the values of ζ_i s in different cliques are independent unless the cliques are neighbors, so Z_j and Z_k are uncorrelated unless $|k - j| \leq 1$. It follows that the conditional variance of $\sum Z_j$ is no larger than 3 times the number of cliques that intersect H . \square

Proof of Proposition 2:. We will compare the distributions of $\sum_{i \in H} \zeta_i(\Pi)$ and $\sum_{i \in H} \zeta_i(\Pi_{\mathcal{G}})$ for a completely random permutation Π independent of \mathcal{G} . Recall that the random variables $\zeta_i(\Pi)$ are Bernoulli- $\frac{1}{2}$, and that $\zeta_i(\Pi)$ and $\zeta_j(\Pi)$ are independent unless $|i - j| \leq 1$. Consequently, $\sum_{i \in H} \zeta_i(\Pi)$ has expectation $|H|/2$ and variance $O(|H|)$, and hence by Chebyshev's inequality

$$P \left\{ \sum_{i \in H} \zeta_i(\Pi) \geq |H|/2 + |H|^{\frac{1}{2} + \frac{\varepsilon}{4}} \right\} \rightarrow 0$$

as $N \rightarrow \infty$.

Now consider the distribution of $\sum_H \zeta_i(\Pi_{\mathcal{G}})$. The set H may be partitioned into nonoverlapping intervals J_1, J_2, \dots, J_m ; by hypothesis (20), $m = O(|H|^{1/2})$. Consequently, by hypotheses (19) and (21) and Lemmas 5–6, with probability approaching 1 as $N \rightarrow \infty$,

$$E \left(\sum_{i \in H} \zeta_i(\Pi_{\mathcal{G}}) \middle| \mathcal{G} \right) \geq \frac{1}{2}|H| + |H|^{\frac{1}{2} + \frac{\varepsilon}{2}}$$

and

$$\text{var} \left(\sum_{i \in H} \zeta_i(\Pi_{\mathcal{G}}) \middle| \mathcal{G} \right) = O(|H|).$$

Thus, by Chebyshev's inequality,

$$P \left\{ \sum_{i \in H} \zeta_i(\Pi_{\mathcal{G}}) \leq |H|/2 + |H|^{\frac{1}{2} + \frac{\varepsilon}{4}} \right\} \rightarrow 0$$

as $N \rightarrow \infty$. This proves that the total variation distance between the distributions of $\sum_{i \in H} \zeta_i(\Pi_{\mathcal{G}})$ and $\sum_{i \in H} \zeta_i(\Pi)$ converges to 1 as $N \rightarrow \infty$, so $d_N(n) \rightarrow 1$.

3.3. Existence of Cold Spots for $C < C_p$. Proposition 2 implies the convergence (17) holds provided there are predictable “cold spots” $H = H_N$ of the deck satisfying relations (19)–(21) as $N \rightarrow \infty$. In this section we will show that such cold spots exist for p near $\frac{1}{2}$ and $n = \lfloor C \log N \rfloor$ for a constant $C \in (0, C_p)$ with C close to C_p .

Define m to be the (least) integer nearest to the solution of

$$(p^{p\theta} q^{q\theta})^m = N^{-\frac{1}{2} + \delta}$$

for a small constant $\delta > 0$ to be specified later. Here $\theta = \theta_p$ where θ_p is defined by equation (5). Observe that m is the (least) integer nearest to the solution m' of

$$(24) \quad m' I(p, p_\theta) = \left(\frac{1}{2} - \delta \right) \log N.$$

Observe that $m < n$, because $p > \frac{1}{2}$ and $p^n < 1/N$, by relation (18). Loosely speaking, the cold spots will be those intervals of the deck in which cards have m -orbits with relative frequency of 1s approximately p_θ . Unfortunately, the precise definition of $H = H_N$ must be somewhat more complicated, because Proposition 2 requires that the sets H be nonrandom.

For each sequence $x = x_1 x_2 \dots x_m \in \Omega_m$, set $t_x = P\{\xi_1 \xi_2 \dots \xi_m \leq x_1 x_2 \dots x_m\}$, where ξ_1, ξ_2, \dots are i.i.d. Bernoulli- p and \leq is the lexicographic order. The numbers t_x , for $x \in \Omega_m$, partition the unit interval into subintervals $J_x = [t_x, t_x + \lambda(x))$ of varying lengths $\lambda(x)$. These intervals are indexed so that if U is uniformly distributed on $[0, 1]$ and $x(U)$ is the index of the interval J_x containing U then the entries x_i of $x(U)$ are i.i.d. Bernoulli- p . Consequently, the length $\lambda(x)$ of the interval J_x coincides with the likelihood $\lambda(x) = \prod_{i=1}^m p^{x_i} q^{1-x_i}$. Fix $\varepsilon > 0$ (small), and define

$$H = H_N = \{i \in [N] : i/N \in \bigcup_{x \in \Omega_m^\varepsilon} J_x\}$$

where

$$\Omega_m^\varepsilon = \Omega_m^\varepsilon(p_\theta)$$

is the set of all m -orbits x for which the relative frequency of 1s is within ε of p_θ . Note that the definition involves the as-yet unspecified constants $\delta > 0$ and $\varepsilon > 0$.

The next two propositions state that for suitably chosen constants $\delta, \varepsilon > 0$, the hypotheses of Proposition 2 are satisfied provided that p is sufficiently near $\frac{1}{2}$ and C is close to C_p . Consequently, their proofs will complete the proof of the assertion (17).

Proposition 3. *For sufficiently small $\varepsilon > 0$, $|H| \rightarrow \infty$ and $|\partial H| = O(|H|^{\frac{1}{2}})$ as $N \rightarrow \infty$.*

Proof. The set H consists of all the integers contained in the nonoverlapping intervals NJ_x , where $x \in \Omega_m^\varepsilon$. Hence, the boundary ∂H has cardinality no larger than $2|\Omega_m^\varepsilon|$. For each $x \in \Omega_m^\varepsilon$, the relative frequency of 1s in x is within ε of p_θ , so the likelihood $\lambda(x)$ (which is also the length of J_x) satisfies

$$\lambda(x) \geq p^{mp_\theta - m\varepsilon} q^{mq_\theta + m\varepsilon} \asymp N^{-\frac{1}{2} + \delta} (p/q)^{-m\varepsilon}.$$

Consequently, the length of each interval NJ_x contained in H is at least a positive multiple (independent of N) of $N^{\frac{1}{2} + \delta} (p/q)^{-m\varepsilon}$. For sufficiently small $\varepsilon > 0$, this is at least $N^{(1+\delta)/2}$. Thus, $|H| \rightarrow \infty$ as $N \rightarrow \infty$. Since $N \geq |H| \geq |\Omega_m^\varepsilon| N^{(1+\delta)/2}$, the cardinality of Ω_m^ε cannot be larger than $N^{(1-\delta)/2}$; it follows that the cardinality of ∂H cannot be larger than $2N^{(1-\delta)/2}$. Since each interval NJ_x contained in H has length at least $N^{(1+\delta)/2}$, this implies that $|\partial H| = o(|H|^{\frac{1}{2}})$. \square

Proposition 4. *Assume that $p - \frac{1}{2}$ is small. For sufficiently small $\delta, \varepsilon > 0$, there exists $\alpha > 0$ such that*

$$(25) \quad \lim_{N \rightarrow \infty} P\{|\mathcal{K} \cap H| \geq |H|^{\frac{1}{2} + \alpha}\} = 1.$$

The remainder of this section is devoted to the proof of this proposition. Since $\mathcal{K} \cap H$ is difficult to deal with directly, we will begin by finding a random set H' , defined solely in terms of orbit properties, for which the size of the intersection $\mathcal{K} \cap H'$ is more easily estimated, and such that $H' \subset H$ with probability approaching 1 as $N \rightarrow \infty$. Fix $0 < \varepsilon' < \varepsilon$; define

$$H' = \{i : x(i) \in \Omega_m^{\varepsilon'}(p_\theta) \cap \Gamma\}$$

where $x(i)$ denotes the m -orbit of the i th card and Γ is the set of all finite 0-1 sequences whose last 2 entries are, in order, 01.

Lemma 7. *Fix $\delta > 0$. For $0 < \varepsilon' < \varepsilon$ and $\varepsilon > 0$ sufficiently small,*

$$\lim_{N \rightarrow \infty} P\{H' \subset H\} = 1.$$

Proof. Recall that the assignment of m -orbits to cards may be made as follows: Start with N independent Bernoulli- p sequences, and attach the labels $i = 1, 2, \dots, N$ to the sequences so that the natural order on the set of labels i coincides with the lexicographic order on the attached sequences; then discard all but the first m entries of each sequence. A stochastically equivalent way to make the assignment is as follows: Drop N points independently from the uniform distribution on $[0, 1]$, and label the points $U_{(i)}$ in accordance with their relative order in $[0, 1]$ (thus, $U_{(1)}, U_{(2)}, \dots, U_{(N)}$ are the order statistics of a random sample of N uniforms). Then replace each $U_{(i)}$ by the (unique) sequence $x = x(i) \in \Omega_m$ such that $U_{(i)} \in J_x$. That these constructions are equivalent follows from the definition of the intervals J_x .

Let $x \in \Omega_m^{\varepsilon'}(p_\theta) \cap \Gamma$. Since $x \in \Gamma$, its last two entries are 01. Consequently, the sequences $x^-, x^+ \in \Omega_m$ immediately to the left and right of x in the lexicographic order are the sequences obtained from x by replacing the last two entries 01 by 00 and 10, respectively. These replacements change the relative frequency of 1s by at most $1/N$ which, for sufficiently large N , is less than $\varepsilon - \varepsilon'$; hence, $x^-, x^+ \in \Omega_m^\varepsilon$. Thus, if $i \in H'$, the intervals $NJ_{x(i)}, NJ_{x(i)^-}$, and $NJ_{x(i)^+}$ are all contained in H .

Let F_N be the empirical distribution function of the sample U_1, U_2, \dots, U_N . The construction of the sequences $x(i)$ forces the inequalities

$$F_N(t_{x(i)^-}) \leq \frac{i}{N} \leq F_N(t_{x(i)} + \lambda(x)).$$

Now let $i \in H'$. By the last paragraph, each of the intervals $NJ_{x(i)}, NJ_{x(i)^-}$ and $NJ_{x(i)^+}$ is contained in H . By the proof of Proposition 3, if $\varepsilon > 0$ is sufficiently small then each interval J_x such that $NJ_x \subset H$ has length at least $|J_x| \geq N^{(-1+\delta)/2}$. Thus, if $i \in H'$, either i/N is in one of the three intervals $J_{x(i)}, J_{x(i)^-}, J_{x(i)^+}$, and hence $i \in H$, or

$$\max_{t \in [0, 1]} \sqrt{N} |F_N(t) - t| \geq N^{\delta/2}.$$

But the Kolmogorov-Smirnov theorem implies that

$$\lim_{N \rightarrow \infty} P\{\max_{t \in [0, 1]} \sqrt{N} |F_N(t) - t| \geq N^{\delta/2}\} = 0.$$

□

Lemma 8. Define τ to be the number of triplets, that is, the number of indices $i \in [N]$ such that cards $i, i+1$, and $i+2$ all have the same n -orbit. Similarly, define κ to be the number of quadruplets, that is, the number of indices $i \in [N]$ such that cards $i, i+1, i+2$, and $i+3$ all have the same n -orbit. For each $\rho > 0$, if p is sufficiently near $\frac{1}{2}$ and if $C_p - C$ is sufficiently small, then as $N \rightarrow \infty$,

$$(26) \quad E\tau = O(N^\rho) \quad \text{and}$$

$$(27) \quad E\kappa = O(N^{-\frac{1}{2}+\rho}) = o(1).$$

Proof. The (unordered) set of n -orbits has the same distribution as a set of N independent, identically distributed Bernoulli- p sequences of length n . There are $\binom{N}{3}$ possibilities for triplets; for each, the probability that their n -orbits coincide is $(p^3 + q^3)^n$. Similarly, there are $\binom{N}{4}$ possibilities for quadruplets; and for each possibility the probability that their n -orbits coincide is $(p^4 + q^4)^n$. Consequently,

$$E\tau = \binom{N}{3} (p^3 + q^3)^n \leq N^3 \exp\{\psi(3)C \log N\} \quad \text{and}$$

$$E\kappa = \binom{N}{4} (p^4 + q^4)^n \leq N^4 \exp\{\psi(4)C \log N\}.$$

As $p \rightarrow \frac{1}{2}$,

$$\begin{aligned} \psi(3) &\longrightarrow -2 \log 2; \\ \psi(4) &\longrightarrow -3 \log 2; \quad \text{and} \\ C_p &\longrightarrow 3/(2 \log 2). \end{aligned}$$

Consequently, if $p - \frac{1}{2}$ is small and $C_p - C$ is small, then relations (26) and (27) will hold as $N \rightarrow \infty$. □

Lemma 9. Assume that p is near $\frac{1}{2}$. Then for every $\delta' \in (0, \delta)$,

$$(28) \quad \lim_{N \rightarrow \infty} P\{|\mathcal{K} \cap H'| \geq N^{1+2\delta'} \exp\{mH(p_\theta) + (n-m)\psi(2)\}\} = 1.$$

Note: Here $H(p_\theta)$ denotes the Shannon entropy function evaluated at p_θ .

Proof. Recall that $i \in \mathcal{K}$ if and only if $(i, i+1) \in \mathcal{E}$, i.e., if and only if the cards i and $i+1$ have the same n -orbit. In addition, $i \in H'$ if and only if its m -orbit is an element of $\Omega_m^{\varepsilon'} = \Omega_m^{\varepsilon'}(p_\theta)$. Consequently, the cardinality of $\mathcal{K} \cap H'$ is bounded below by the number of distinct n -orbits $x = x_1 x_2 \dots x_n$ such that (i) at least two cards have n -orbit x ; and (ii) $x_1 x_2 \dots x_m \in \Omega_m^{\varepsilon'}$. Thus, by the inclusion-exclusion formula,

$$|\mathcal{K} \cap H'| \geq \sum_{x \in \Omega_m^{\varepsilon'}} (\xi_2(x) - \xi_3(x) - \xi_4(x) - \dots)$$

where $\xi_k(x)$ is the number of k -sets of cards with the same n -orbit y and $y_1 y_2 \dots y_m = x$. Now

$$\sum_{x \in \Omega_m^{\varepsilon'}} \xi_3(x) \leq \tau \quad \text{and} \quad \sum_{x \in \Omega_m^{\varepsilon'}} \xi_k(x) \leq \kappa \quad \forall k \geq 4,$$

where τ and κ are the number of triplets and quadruplets, respectively. By Lemma 8, for each $\rho > 0$ there exists $\beta = \beta_\rho > 0$ such that if $p - \frac{1}{2} < \beta$ and $C_p - C < \beta$ then $E\tau < N^\rho$ and $E\kappa \leq N^{-(1/2)+\rho}$ for all sufficiently large N . Hence, by the Markov inequality,

$$\lim_{N \rightarrow \infty} P\{\tau > N^{2\rho}\} = 0 \quad \text{and} \quad \lim_{N \rightarrow \infty} P\{\kappa \geq 1\} = 0.$$

We will prove below that $Y := \sum_{x \in \Omega_m^{\varepsilon'}} \xi_2(x)$ is, with probability approaching 1 as $N \rightarrow \infty$, much larger than $N^{2\rho}$. Therefore, to prove Lemma 9 it suffices to prove that (28) holds with $|\mathcal{K} \cap H'|$ replaced by Y .

We shall estimate the expectation and variance of the random variable Y by appealing to Corollary 4 of section 2.4. First,

$$EY = E \sum_{x \in \Omega_m^{\varepsilon'}} \xi_2(x) = \binom{N}{2} (p^2 + q^2)^{n-m} \sum_{x \in \Omega_m^{\varepsilon'}} \lambda(x)^2.$$

By Corollary 4, there exist constants $\gamma = \gamma(\varepsilon')$ satisfying $\gamma \rightarrow 0$ as $\varepsilon' \rightarrow 0$ such that for all sufficiently large N ,

$$\sum_{x \in \Omega_m^{\varepsilon'}} \lambda(x)^2 \geq \exp\{mH(p_\theta) - 2mI(p, p_\theta) - m\gamma\};$$

hence, by equation (24), if $\varepsilon' > 0$ is sufficiently small then

$$\begin{aligned} EY &\geq \frac{1}{2} \left(1 - \frac{1}{N}\right) N^{1+2\delta} \exp\{mH(p_\theta) - m\gamma + (n-m)\psi(2)\} \\ (29) \quad &\geq N^{1+2\delta''} \exp\{mH(p_\theta) + (n-m)\psi(2)\} \end{aligned}$$

for some $0 < \delta' < \delta'' < \delta$. Observe that, if p is near $\frac{1}{2}$ and C is near C_p , this is at least $N^{(1/2)-\rho}$ for arbitrarily small $\rho > 0$, and so is of larger order of magnitude than $E\tau$.

The variance of Y is also easily estimated. Write

$$Y = \sum_{x \in \Omega_m^{\varepsilon'}} \xi_2(x) = \sum_{i, j \in [N]: i < j} Y_{ij},$$

where Y_{ij} is the indicator of the event that the n -orbits x^i and x^j coincide and their common truncations fall in $\Omega_m^{\varepsilon'}$ (here the n -orbits are taken to be randomly ordered). The random variables Y_{ij} are identically distributed Bernoullis. If i, j, i', j' are distinct indices then Y_{ij} and $Y_{i'j'}$ are independent; if i, i', j are distinct indices, Y_{ij} and $Y_{i'j}$ are not independent, but $\sum_{i < i' < j} Y_{ij} Y_{i'j} = \tau$, which has expectation $O(N^\rho)$. Consequently,

$$(30) \quad \text{var}(Y) \leq E \sum_{x \in \Omega_m^{\varepsilon'}} \xi_2(x) + O(N^\rho).$$

The result now follows from relations (29)–(30), by Chebyshev's inequality. \square

Proof of Proposition 4. Recall that m is the (least) integer nearest the solution m' of equation (24). Recall that H consists of the nonoverlapping intervals NJ_x , where $x \in \Omega_m^\varepsilon$. Since the length of J_x is $\lambda(x)$, Corollary 4 implies that, for some $\gamma = \gamma(\varepsilon) > 0$ satisfying $\gamma \rightarrow 0$ as $\varepsilon \rightarrow 0$,

$$|H| = N \sum_{x \in \Omega_m^\varepsilon} \lambda(x) \leq N e^{m'H(p_\theta) - m'I(p, p_\theta) + m'\gamma} = N^{\frac{1}{2} + \delta} e^{m'H(p_\theta) + m'\gamma}.$$

Therefore, by Lemmas 7 and 9, it suffices to prove that for some $\alpha > 0$ and some $0 < \delta' < \delta$,

$$\left(N^{\frac{1}{2}+\delta} \exp\{mH(p_\theta) + m\gamma\}\right)^{\frac{1}{2}+\alpha} \leq N^{1+2\delta'} \exp\{mH(p_\theta) + (n-m)\psi(2)\}.$$

Each side of this inequality may be expressed as N raised to a power, using equations (16) and (24) relating n, m , and N . It suffices to show that the power on the left side is less than the power on the right side. Since γ, δ , and α may be taken arbitrarily small, they may be set equal to zero before comparing the powers. Thus, it suffices to prove that

$$1 + \frac{H(p_\theta)}{I(p, p_\theta)} < 4 + 2\frac{H(p_\theta)}{I(p, p_\theta)} + 2\left(2C - \frac{1}{I(p, p_\theta)}\right)\psi(2).$$

Now recall (equation (15)) that $H(p_\theta) = \theta I(p, p_\theta) + \psi(\theta)$; hence, the last inequality reduces to

$$3 + \theta + \frac{\psi(\theta)}{I(p, p_\theta)} + 2\psi(2)\left(2C - \frac{1}{I(p, p_\theta)}\right) > 0.$$

But $\psi(\theta) = 2\psi(2)$, so this is equivalent to

$$3 + \theta + 4\psi(2)C > 0.$$

Since $0 = 3 + \theta + 4\psi(2)C_p$, by definition of C_p , and since $C < C_p$ and $\psi(2) < 0$, the inequality $3 + \theta + 4\psi(2)C > 0$ is in fact true.

4. A HEURISTIC ARGUMENT

In this section we give a heuristic argument in favor of Conjecture 1. This argument is based on a rough analysis of the Radon–Nikodym derivative dQ/dP , where P and Q are, respectively, the uniform distribution and the distribution of the deck after n repetitions of the p -shuffle.

Recall (Lemma 3) that a version $\Pi_{\mathcal{G}}$ of the random permutation induced by n repetitions of the p -shuffle may be constructed as follows: Start with a completely random (uniformly distributed) permutation Π . Build a random graph \mathcal{G} from an $(N \times n)$ Bernoulli matrix X , independent of Π , by putting edges between neighboring indices $i, i+1$ such that the i th and $(i+1)$ th largest (lexicographically) rows of X are the same. Then use the random graph \mathcal{G} to modify Π by ordering assignments in \mathcal{G} -cliques. Let Q and P be the distributions of the random permutations $\Pi_{\mathcal{G}}$ and Π , respectively.

Lemma 8 above implies that if $p - 1/2$ and $C - C_p$ are small (here $n = \lfloor C \log N \rfloor$) then (i) the expected number of “triplets” (vertices i such that the graph \mathcal{G} contains edges connecting i to $i+1$ and $i+1$ to $i+2$) is of order no more than N^ρ , where $\rho < 1/2$; and (ii) the expected number of “quadruplets” is $o(1)$. If \mathcal{G} had neither triplets nor quadruplets, then for any permutation π , the only information about which of the distributions P, Q might have “generated” π in the values $\pi(i)$ and $\pi(i+1)$ would be in the value of $\zeta_i(\pi)$, the indicator of the event $\pi(i) < \pi(i+1)$. Under P , this random variable is Bernoulli- $1/2$, while under Q , it is Bernoulli- q_i . Here $q_i = (1/2)(1 + \rho_i)$, and

$$\rho_i = \rho_i^{N,n} = P(K_i),$$

where $K_i = K_i^{N,n}$ is the event that the orbits of cards i and $i+1$ coincide up to time n . If the information in the random variables ζ_i were independent, then the Radon–Nikodym

derivative (dQ/dP) would be the product of the likelihood ratios for the Bernoulli random variables ζ_i :

$$(31) \quad \frac{dQ}{dP}(\pi) = \prod_{i=1}^{N-1} (1 + \rho_i)^{\zeta_i} (1 - \rho_i)^{1-\zeta_i}.$$

If n were sufficiently large that all of the coincidence probabilities ρ_i were small, then (31) could be rewritten as

$$(32) \quad \frac{dQ}{dP}(\pi) = 1 + \sum_{i=1}^{N-1} \xi_i \rho_i + \cdots,$$

where $\xi_i = 2\zeta_i - 1$ and \cdots indicates higher order terms. Now the random variables ξ_1, ξ_3, \dots are independent under P , as are the random variables ξ_2, ξ_4, \dots , and each ξ_i has mean zero and variance 1 under P . Consequently, if $n = \lfloor C \log N \rfloor$ for some $C > C_p$, then Proposition 5 below would imply that, under P , as $N \rightarrow \infty$,

$$(33) \quad \frac{dQ}{dP} \xrightarrow{P} 1,$$

which would, in turn, imply the truth of Conjecture 1.

Proposition 5. *Let $n = \lfloor C \log N \rfloor$. Then*

$$(34) \quad C > C_p \implies \lim_{N \rightarrow \infty} \sum_{i=1}^N \rho_i^2 = 0;$$

$$(35) \quad C < C_p \implies \lim_{N \rightarrow \infty} \sum_{i=1}^N \rho_i^2 = \infty.$$

We shall omit the proof of this result, since we have not been able to make the rest of the argument rigorous. It seems likely to us that, although the expression (31) for the Radon–Nikodym derivative dQ/dP is not exact, it is nevertheless close enough to the true value that the remainder of the argument remains valid.

Acknowledgments. Thanks to PERSI DIACONIS for his encouragement, and to JIM FILL for his incredibly thorough reading of the paper and numerous helpful suggestions.

REFERENCES

- [1] D. Aldous. Random walks on finite groups and rapidly mixing Markov chains. *Seminar on probability XVII*, 243–297, Lecture Notes in Math., 986, Springer, Berlin–New York, 1983.
- [2] D. Aldous and P. Diaconis. Strong uniform times and finite random walks. *Adv. in Appl. Math.* 8 (1987), no. 1, 69–97.
- [3] D. Bayer and P. Diaconis. Trailing the dovetail shuffle to its lair. *Ann. Appl. Probab.* 2 (1992), no. 2, 294–313.
- [4] P. Bidigare, P. Hanlon, and D. Rockmore. A combinatorial description of the spectrum of the Tsetlin library and its generalization to hyperplane arrangements. *Duke Math. J.* 99 (1999), 135–174.
- [5] K. Brown and P. Diaconis. Random walks and hyperplane arrangements. *Ann. Probab.* 26 (1998), 1813–1854.
- [6] P. Diaconis. *Group representations in probability and statistics* Institute of Mathematical Statistics Lecture Notes–Monograph Series, 11. Institute of Mathematical Statistics, Hayward, CA, 1988
- [7] P. Diaconis, J. Fill, & J. Pitman. Analysis of top to random shuffles. *Combinatorics, Probability, & Computing* 1 (1992), 135–155.
- [8] J. Fulman. The combinatorics of biased riffle shuffles. *Combinatorica* 18 (1998), 173–184.

DEPARTMENT OF STATISTICS, UNIVERSITY OF CHICAGO, ECKHART HALL, 5734 UNIVERSITY AVENUE
CHICAGO, ILLINOIS 60637

E-mail address: `lalley@galton.uchicago.edu`