from *SIAM News*, Volume 36, Number 1, January/February 2003

Within Every Math Problem, For this Mathematician, Lurks a Card-Shuffling Problem

By Erica Klarreich

Sorting through his mail after winter break last year, Persi Diaconis came upon a letter containing a request that couldn't fail to whet his interest. A large manufacturer of card-shuffling machines for casinos had just designed a new type of shuffler, but wasn't sure whether the decks it spewed out were really random. The company asked Diaconis to visit its Las Vegas headquarters and tell company engineers whether the new machine made the grade.

Diaconis was a smart choice, for he has a degree of intimacy with cards that few mathematicians can claim. He came to mathematics via the life of a professional itinerant magician. At fourteen, he ran away from home to go on the road with Dai Vernon, whom Diaconis calls "the Einstein of magic in the twentieth century." For a decade, the pair scoured the globe for anyone who could teach them something new about a deck of cards. "If we heard there was an Eskimo who had a way of dealing the second card with snowshoes, we'd be off to Alaska," Diaconis says.

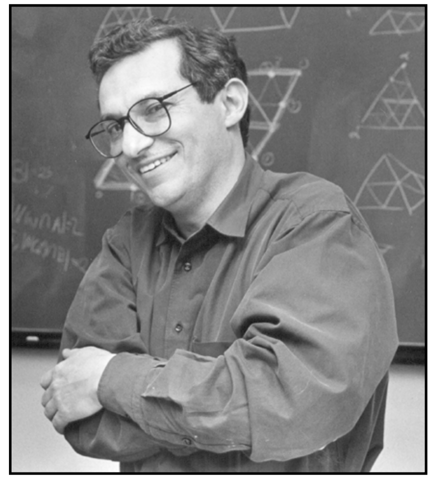

With a Stanford colleague, Persi Diaconis demonstrated mathematically that a Las Vegas-destined card shuffler fell far short of randomizing a deck of cards. But a tough challenge remained: convincing casino executives, "who don't know that they care about my esoteric total variation distance," that the machine was unacceptable for casino play.

Eventually, a meeting with Martin Gardner made Diaconis realize the fascination of mathematics, and after 10 years of this wandering life he returned to school, finishing an undergraduate degree and a PhD within five years. Now a professor in the mathematics and statistics departments at Stanford University, Diaconis has a trick of turning every mathematics problem he comes upon into a question about card shuffling. "If I can manage to translate a problem into something about shuffling, then I have a shot at being able to solve it," he says. Ten years ago, he and David Bayer, now at Columbia University, caught the attention not only of the mathematics community but also of the popular press by proving that a deck is randomized after about seven ordinary "riffle" shuffles, in which the deck is divided into two piles that are then interleaved with each other.

The company's new machine was a "shelf shuffler," which operates by dropping cards onto shelves, one after another, and then gathering up the piles. The machine starts at the bottom of the deck; for each card, a random number generator determines which of ten shelves to put the card on, after which a second random number generator decides whether to put the card on the top or bottom of the pile of cards already on that shelf. When all the cards have been placed, the machine spits out the packets in random order.

A slew of shelf shufflers have come onto the market lately, Diaconis says. But this particular shelf shuffler ended up on the shelf, not in a Las Vegas casino—after Diaconis told the company engineers all the things that were wrong with it.

At first glance, the machine seems to build in an abundance of randomizing features. To understand better just what the machine does, Diaconis and his collaborator Susan Holmes, a statistician at Stanford, created a model to abstract the key ingredients of the machine. They noticed that when all the cards have been distributed on shelves, each shelf holds two sequences of cards. The cards that were placed on top of the pile already on the shelf are in the same relative order in which they appeared in the deck, while the cards that were placed on the bottom of the pile are in reverse order. So the output of a shuffle consists of ten packets, each of which contains one sequence of cards in ascending order and one in descending order. That means you can achieve exactly the same result by randomly assigning to each card a number between 1 and 20, and then gathering all the cards with a given number into a single pile and stacking the piles; before stacking, reverse the order of the cards in piles 2, 4, 6, and so on, but preserve the order of the cards in piles 1, 3, 5, and so on.

This model makes it clear that the last act of the machine—the random rearrangement of the packets—achieves nothing, since it amounts to changing the numbers assigned to the cards, which were chosen randomly anyway. The company's engineers were taken aback but they were also delighted at this discovery, since it meant they could put the capacity that went into the random rearrangement to other uses. "That's why they had gone to a mathematician," Diaconis says. "Sometimes the power of abstract thought defeats your intuition."

But the model showed up even more serious shortcomings of the shelf shuffler, too. Assume that a deck placed in the machine starts out in some canonical order—say, from top to bottom, ace through king of spades, then ace through king of hearts, then diamonds, then clubs. The ace of spades has a 1 in 20 chance of being assigned to pile number 1, and if it is, it must be on the top of the pile, since it is the top card in the deck. So the probability that the ace of spades is still at the top of the deck after the shuffle is 1 in 20, not the 1 in 52 it should be.

It was clear that the shuffler falls far short of randomizing the deck. To figure out how it stacks up against riffle shuffling, Diaconis and Holmes computed the total variation distance between the shuffler's set of outcomes and the uniform distribution of arrangements of 52 cards. For a shuffling method Q, the total variation distance after n shuffles is given by the formula

$$||Q^n - U|| = \max_{A \subset S_n} |Q^n(A) - U(A)|,$$

where S_n is the symmetric group of permutations of a set of n elements, U is the uniform distribution, and $Q^n(A)$ is the probability distribution after n shuffles.

The pair realized that all the necessary ingredients for analyzing the total variation distance of the shelf shuffler were embedded in the riffle-shuffling analysis Diaconis and Bayer had done ten years earlier. At that time, they had studied "type b" shuffles, in which the deck is divided into two piles, and one of them is turned face up before the two are riffled together. Diaconis and Bayer had shown that performing n type b shuffles results in the same total variation distance as performing $n - 1$ ordinary riffle shuffles.

The standard type b shuffle involves only two packets, but it is easy to perform type b shuffles with more than two packets: A six-packet type b shuffle, for instance, would involve turning three of the packets face up, then interleaving all the packets together. A little thought shows that if you ignore the fact that some cards are face up at the end of a type b shuffle, the 10-platform shelf shuffle is the inverse of the 20-packet type b shuffle. A 20-packet type b shuffle, after all, takes 20 packets, reverses the order of half of them (by turning them face up), then interleaves them. A 10-platform shelf shuffle exactly undoes this, grouping into 20 packets cards that had been interleaved, and then reversing the order of half the packets. Since the two kinds of shuffles are inverses, it is not hard to see that they will produce the same total variation.

"A variety of combinatorialists are now thinking of card shuffling as illuminating their areas."

Doing a type b shuffle with four packets is the same as doing two two-packet type b shuffles—both will result in four interleaved packets, two face up and two face down. In the same way, doing a type b shuffle with 2^n packets is the same as doing n two-packet type b shuffles. So doing a type b shuffle with 20 packets is like doing somewhere between four and five two-packet type b shuffles, and by Diaconis and Bayer's analysis, that gives the same total variation distance as doing between three and four ordinary riffle shuffles. The company's shelf shuffler, then, wasn't even as good as four riffle shuffles—nowhere near the seven shuffles needed to randomize a deck.

Having settled to their own satisfaction that the machine didn't work, Diaconis and Holmes were left with the problem of demonstrating it to the satisfaction of their intended audience. "The task was to go back to the casino executives who have a real vested interest in their machines, and who don't have a mathematical background," Diaconis says. "They don't know that they care about my esoteric total variation distance. What they care about is whether it will make a difference in the kind of casino games they use the machine for."

To drive home the point that the machine's output would not be acceptable in casino play, Holmes decided to consider a question that casino operators care about: How many cards could a gambler guess correctly in a deck that has passed through the shelf shuffler? In a deck that has been properly randomized, the gambler has a 1 in 52 chance of guessing the top card correctly; after the top card has been revealed, the gambler has a 1 in 51 chanced of guessing the second card correctly; and so on. Going through the entire deck, a guesser can expect to get about 4.5 cards right.

But the shelf shuffler—with its arrangement of the deck into 20 packets, alternating between ascending and descending order—doesn't put the deck into random order. Using that fact, Holmes came up with a guessing strategy for a shelf-shuffled deck that enables a gambler to guess about 9.5 cards correctly, provided the gambler knows the order of the cards before they go through the shuffler.

Suppose the deck is in the canonical order before it is shuffled. Then the gambler should guess that the top card is the ace of spades, since she has a 1 in 20 chance of being right. If this first guess is correct, she should next guess the two of spades, since it also has a 1 in 20 chance of being in the top packet, and if it is there, it must come right after the ace. If the first guess is incorrect and the top card is not the ace of spades but, say, the six of diamonds, then she should guess the seven of diamonds next, for the same reason. The gambler should continue moving her guesses upward in this way—if the three of clubs shows up, for instance, she should next guess the four of clubs—until a card is turned up that is lower than the card that preceded it. That signals the start of the next packet, which is in descending rather than ascending order. Now the gambler should reverse her strategy: If, say, the eight of hearts turns up, she should guess the seven of hearts (unless it has already been turned up, in which case she should guess the six of hearts). She should keep guessing downward until a card turns up that is higher than the one before it; that is the signal to start guessing upward again. Holmes's students proved that this switching strategy is the optimal way to guess through the deck.

Holmes's strategy convinced the company's engineers that the machine was absolutely unacceptable in its present form. How they will alter the machine is something they are not letting on, even to Diaconis and Holmes. But the mathematical analysis shows that there is at least one quick fix. Since the machine randomizes the deck about as well as three or four riffle shuffles, running a deck through the machine twice will produce a total variation distance roughly equal to that of seven riffle shuffles, which Diaconis and Bayer's earlier analysis showed to be acceptable for casino play.

The questions raised by the shelf shuffler are continuing to give rise to new mathematics. While analyzing the machine, Diaconis came up with a closed formula for the expected number of fixed points of a shuffle—cards that are still in their original position after the shuffle. "I was surprised that such a complicated randomizing device would yield such a formula," Diaconis says. He mentioned the formula to Jason Fulman, a professor at the University of Pittsburgh who was then a postdoc at Stanford; Fulman went on to figure out formulas for the likelihood of more complicated permutations arising as a result of the shuffle.

"He derived these beautiful formulas using the most esoteric Lie theory," Diaconis says. "It's math that has nothing to do with shuffling cards but it was there, and he knew about it." Fulman's formulas created an unexpected link between shelf shuffling and symmetric function theory, which looks at functions that are unchanged when some of the variables are permuted. "A variety of combinatorialists are now thinking of card shuffling as illuminating their areas," Diaconis says.

For Diaconis, the shelf-shuffling work is a vindication of his strategy of working on any problem that sparks his curiosity. "The only reason we had studied the type b shuffles in the first place was that they were beautiful," he says. "We didn't know they would later turn out to be exactly what was needed to understand this application."

Erica Klarreich is a freelance writer based in Berkeley, California.