SUMMARY OF CARD SHUFFLING PAPER

BRIAN CITRO

In Brad Mann's paper entitled "How Many Times Should You Shuffle a Deck of Cards?" (which can be found at `http://www.dartmouth.edu/~chance/teaching_aids/Mann.pdf`), the author attempts to determine the proper amount of shuffling required to adequately randomize a standard deck of cards. Mann begins by defining a shuffle as a probability density on the total number of possible permutations of the cards, with each permutation given a fixed probability of occurring. Thus, a repeated shuffle is equal to a random walk of all possible permutations of the cards, and the overall probability density after a fixed number of steps is equal to the convolution of each shuffle's density.

Mann next considers how people actually shuffle cards. He develops a model of a riffle shuffle, in which a person cuts the deck into two piles, with the number of cards cut corresponding to a binomial density. The cards are then interleaved, with each possible interleaving being equally likely. After building this model, Mann tries to determine how many convolutions of riffle shuffles are necessary to randomize the deck. In order to do this, he first must define what it means for the deck to be randomized. Mann states that in the ideal random deck each ordering of cards would be equally likely, with each ordering occurring with probability $1/52!$. However, there is no number of riffle shuffles that could be performed that could achieve true randomness. Therefore, a metric for measuring how close to true randomness a density is must be used, which in this case is the variation distance between the riffle shuffle density and the uniform density.

In order to actually determine the density of repeated riffle shuffles, Mann introduces the concept of a rising sequence, which is the maximal consecutively increasing subsequence of a permutation. Using this idea, Mann is able to calculate the variation distance after k shuffles. He concludes that the graph of his results begins to decline at around 5 shuffles, and by 11 shuffles the result is reasonably close to 0. This explains the previously established notion that seven shuffles is enough for a deck of cards. This research has interesting implications for casinos, which rely on randomness to ensure that the house maintains an advantage.