

Online card games*

Sam Spiro[†]

Abstract

Consider the following one player game. A deck containing m copies of n different card types is shuffled uniformly at random. Each round the player tries to guess the next card in the deck, and then the card is revealed and discarded. It was shown by Diaconis, Graham, He, and Spiro that if m is fixed, then the maximum expected number of correct guesses that the player can achieve is asymptotic to $H_m \log n$, where H_m is the m th harmonic number.

In this paper we consider an adversarial version of this game where a second player shuffles the deck according to some (possibly non-uniform) distribution. We prove that a certain greedy strategy for the shuffler is the unique optimal strategy in this game, and that the guesser can achieve at most $\log n$ expected correct guesses asymptotically for fixed m against this greedy strategy.

Keywords: game theory; card shuffling; discrete probability.

MSC2020 subject classifications: 60C05.

Submitted to EJP on July 17, 2021, final version accepted on March 23, 2022.

Supersedes arXiv:2106.11866v2.

1 Introduction

Consider the following one player game played by a player Guesser. A deck with n different card types with each card type appearing m times is shuffled uniformly at random. Each round, Guesser guesses what the card type of the next card in the deck is. After the guess is made, the true identity of the card is revealed and then discarded from the deck. This process continues until the deck is depleted, at which point the game ends and Guesser is given a point for each correct guess they made during the game.

This game is called the *complete feedback model*, and it was motivated by several real world problems related to clinical trials [4, 10] and to extrasensory perception experiments [6]. We refer the interested reader to [8] for more information on the history of this model, as well as to variants of the model which involve different levels of feedback.

*This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1650112.

[†]University of California San Diego, United States of America E-mail: sspiro@ucsd.edu

A natural question to ask regarding the complete feedback model is how many points Guesser can expect to guess if they try to maximize or minimize their score. To this end, if \mathcal{G} is a strategy in the complete feedback model¹ and if π is a way to shuffle a deck which consists of m copies of n different card types, we define $C_{m,n}(\mathcal{G}, \pi)$ to be the number of correct guesses made by Guesser in the complete feedback model if they use the strategy \mathcal{G} and if the deck is shuffled according to π . Building on work of Diaconis and Graham [6], the following was proven by Diaconis, Graham, He, and Spiro [7] and by He and Ottolini [14]. Here and throughout $\log n$ denotes the natural logarithm and $\Gamma(x)$ denotes the gamma function.

Theorem 1.1 ([6, 7, 14]). *For any fixed m and n tending towards infinity, we have*

$$\begin{aligned} \max_{\mathcal{G}} \mathbb{E}[C_{m,n}(\mathcal{G}, \pi)] &\sim H_m \log n, \\ \min_{\mathcal{G}} \mathbb{E}[C_{m,n}(\mathcal{G}, \pi)] &\sim \Gamma\left(1 + \frac{1}{m}\right) n^{-1/m}, \end{aligned}$$

where the maximum and minimums run over all strategies \mathcal{G} , the shuffling π is chosen uniformly at random, and $H_m := \sum_{i=1}^m i^{-1}$ is the m -th harmonic number. Moreover, \mathcal{G} achieves this maximum/minimum if it guesses a most/least likely card each round.

We note that the regime where n is fixed and m tends to infinity is studied in [6], and in this case both the maximum and minimum scores are asymptotic to m .

The focus of this paper is in studying the complete feedback model with a deck of cards that is not shuffled uniformly at random, which is a reasonable assumption to make when considering real life applications. There is a vast literature concerning different models for shuffling a deck of cards, and we briefly summarize some of these models below. For the following discussion, we restrict ourselves to decks consisting of n distinct cards (so $m = 1$ in our language).

One popular way to shuffle cards (both in real life and mathematically) is the *riffle shuffle*, which informally involves splitting the deck roughly in half and then interleaving the two halves together. A mathematical framework for modeling riffle shuffles was established by Gilbert, Shannon, and Reeds [13], and it was shown by Bayer and Diaconis [3] that roughly $\frac{3}{2} \log_2 n$ riffle shuffles suffices to produce a deck which is “close” to being shuffled uniformly. The $n = 52$ case of their analysis has led to the maxim that seven riffle shuffles suffice to make a standard deck of playing cards close to uniformly shuffled.

A simpler model of shuffling cards is the *top to random* shuffle, which involves taking the topmost card of the deck and inserting it uniformly at random into the deck. It was shown by Aldous and Diaconis [2] that it takes about $n \log n$ top to random shuffles to get a deck which is close to uniformly shuffled. There are many other kinds of shuffling models that can be considered, such as random transpositions [9] and overhand shuffles [15].

Given these various models for how to shuffle a deck of cards, it is reasonable to ask how well a player can do in the complete feedback model if the deck is shuffled in some non-uniform way. Specifically, Bayer and Diaconis [3] asked what the optimal score in the complete feedback model is when the deck is shuffled using k riffle shuffles. The case of $k = 1$ was essentially solved by Liu [18] and later refined by Krityakierne and Thanatipanonda [17]. Ciucu [5] established similar results in the no feedback model (i.e. where the Guesser is told nothing after each successive guess) after the deck has been riffle shuffled once. Further work in the no feedback model was done by Pehlivan [19] in the setting of top to random shuffles.

¹We formally define what a “strategy” is in Section 2. Intuitively, it is a function which takes in the information which is currently known to Guesser and outputs a (possibly random) card type to guess.

In this paper we consider a somewhat more general situation by allowing any distribution for shuffling of the cards to be used in the complete feedback model. More precisely, we consider the following two player game played by Guesser and Shuffler. The game starts with an initial deck of cards. Each round Shuffler chooses a card which is in the deck², and simultaneously Guesser guesses the card type of the selected card. The card is then revealed and discarded. The game continues in this way until the deck is completed. We call this the *Guesser-Shuffler game*. We define $C_{m,n}(\mathcal{G}, \mathcal{S})$ to be the number of correct guesses made by Guesser in this game if the deck starts with m copies of n different card types and if Guesser and Shuffler use the strategies \mathcal{G} and \mathcal{S} , respectively.

For example, Shuffler could use the strategy “shuffle the deck uniformly at random at the start of the game, then sequentially pick the top card from the deck each round.” Under this strategy the game reduces to the complete feedback model, and we know from Theorem 1.1 that Guesser can play so that they get $H_m \log n$ points in expectation. The central question we wish to answer is: can Shuffler significantly decrease this quantity if they use a different strategy?

To this end, we say that a strategy \mathcal{S} for Shuffler is a *minimizing Shuffler strategy* if for all m, n we have

$$\max_{\mathcal{G}} \mathbb{E}[C_{m,n}(\mathcal{G}, \mathcal{S})] = \min_{\mathcal{S}'} \max_{\mathcal{G}} \mathbb{E}[C_{m,n}(\mathcal{G}, \mathcal{S}')],$$

where here the maximums and minimums range over all strategies for Guesser and Shuffler, respectively. We similarly define what it means for a strategy to be a *maximizing Shuffler strategy*. To state the main result of this paper, we define the following.

Definition 1.2. *The greedy strategy \mathcal{S}_g for Shuffler is defined as follows. Each round, if the deck has s different card types remaining in the deck, then Shuffler selects each card type with probability proportional to s^{-1} .*

For example, if at some point the deck consisted of 100 copies of one card type and 1 copy of another card type, then \mathcal{S}_g would select each of these card types with probability $\frac{1}{2}$. One can easily show that the greedy strategy minimizes the probability that Guesser correctly guesses a card type in any given round, but it is far from obvious that it is a minimizing Shuffler strategy. Indeed, in our example, if Shuffler chooses the card type with multiplicity 1, then Guesser can guarantee 100 correct guesses from later rounds. Thus it might be the case that Shuffler should choose this low multiplicity card type with probability smaller than $\frac{1}{2}$. However, this turns out not to be the case.

Theorem 1.3. *The greedy strategy \mathcal{S}_g is both a minimizing and maximizing Shuffler strategy for the Guesser-Shuffler game. Moreover, it is the unique minimizing Shuffler strategy.*

We note that there is no unique maximizing Shuffler strategy. Indeed, once a card type i has been exhausted from the deck, Guesser can minimize their score by guessing i every round, and at this point the strategy that Shuffler uses is irrelevant. However, it is possible to show that every maximizing Shuffler strategy must agree with the greedy strategy until a card type runs out of the deck, see Theorem 2.5 for a precise statement of this form.

By using Theorem 1.3 together with results from variants of the coupon collector and birthday problems, we can prove asymptotic bounds on how many correct guesses are

²In this formulation of the game, Shuffler chooses each card in an online fashion, possibly based on what Guesser has done in previous rounds. Alternatively one could consider a model where Shuffler shuffles the deck at the start of the game and then draws each card according to this shuffling. It is not difficult to show that the optimal expected scores under these two models are the same, but for concreteness we will only consider the online setting.

made in the Guesser-Shuffler game when one player tries to minimize the number of correct guesses and the other tries to maximize it.

Theorem 1.4. *For m fixed and n tending towards infinity, we have*

$$\min_{\mathcal{G}} \max_{\mathcal{S}} \mathbb{E}[C_{m,n}(\mathcal{G}, \mathcal{S})] = \max_{\mathcal{G}} \mathbb{E}[C_{m,n}(\mathcal{G}, \mathcal{S}_g)] \sim \log n, \tag{1.1}$$

$$\max_{\mathcal{S}} \min_{\mathcal{G}} \mathbb{E}[C_{m,n}(\mathcal{G}, \mathcal{S})] = \min_{\mathcal{G}} \mathbb{E}[C_{m,n}(\mathcal{G}, \mathcal{S}_g)] \sim \Gamma\left(1 + \frac{1}{m}\right) (m!)^{1/m} \cdot n^{-1/m}. \tag{1.2}$$

In particular, comparing Theorem 1.4 with Theorem 1.1, we see that Shuffler can significantly decrease or increase the expected number of correct guesses if they use the greedy strategy as opposed to shuffling the deck uniformly at random.

To give some intuition for the bound in (1.1), let \mathcal{G}^+ be the Guesser strategy which uniformly at random chooses a card type which has the maximum number of copies left in the deck. It is not hard to show that the expected score under this strategy is at least the n th harmonic number $H_n \sim \log n$ regardless of how the deck is shuffled. Thus (1.1) can be interpreted as saying that no strategy for Guesser can asymptotically outperform \mathcal{G}^+ when the deck is shuffled adversarially. This is in contrast to having the deck shuffled uniformly at random, since in that setting Theorem 1.1 shows that \mathcal{G}^+ achieves significantly higher scores for larger values of m .

Lastly, we consider a variant of the Guesser-Shuffler game that we call the *Restricted Matching Pennies game*. This is a two player game which starts with both players being given an identical deck of cards. Each round the two players simultaneously select a card from their deck, reveal their cards, and then discard them. This continues until the two decks are depleted. We define $M_{m,n}(\mathcal{A}, \mathcal{B})$ to be the number of times that the selected cards matched during a round of the game if the deck starts with m copies of n different card types and if the two players use strategies \mathcal{A} and \mathcal{B} .

One can view this game as a symmetric version of the Guesser-Shuffler game where Guesser is forced to guess each card type exactly m times (with Guesser’s deck recording the remaining guesses that they are allowed to make). This can also be seen as a variant of the classical game “Matching Pennies.” In each round of this game, two players simultaneously reveal a penny from their hand, where one player gets a point if the coins match (i.e. if they are both heads or both tails), and the other gets a point if they do not match. In contrast, the Restricted Matching Pennies game allows n different types objects to be chosen from each round, and it adds the restriction that the number of times each object is chosen in total is fixed.

It turns out that in the Restricted Matching Pennies game that the optimal strategy for both players is to shuffle their decks uniformly at random, which is in sharp contrast to Theorem 1.3.

Proposition 1.5. *For all m, n we have*

$$\max_{\mathcal{A}} \min_{\mathcal{B}} \mathbb{E}[M_{m,n}(\mathcal{A}, \mathcal{B})] = m.$$

Moreover, if \mathcal{U} is the strategy of choosing each card from the remaining deck uniformly at random, then for any strategies \mathcal{A}, \mathcal{B} we have

$$\mathbb{E}[M_{m,n}(\mathcal{A}, \mathcal{U})] = \mathbb{E}[M_{m,n}(\mathcal{U}, \mathcal{B})] = m.$$

The rest of this paper is organized as follows. In Section 2 we prove Theorem 1.3 showing that the greedy strategy is optimal in the Guesser-Shuffler game. We then prove the asymptotic bounds of Theorem 1.4 in Section 3 and the results for the Restricted Matching Pennies game in Section 4. Concluding remarks and some open problems are given in Section 5.

2 Proof of Theorem 1.3

To prove our results, we first formally define some of the terms used in the introduction.

We say that d is an n -vector if it is a vector (d_1, \dots, d_n) of non-negative integers, and we think of d_i as the number of card types of type i left in a deck of cards. We define the support of d to be $\text{supp}(d) = \{i : d_i > 0\}$, and if \mathbf{X} is a random variable we define its support by $\text{supp}(\mathbf{X}) = \{x : \Pr[\mathbf{X} = x] > 0\}$.

We say that \mathcal{G} is a *Guesser strategy* if it is a function which takes in nonzero n -vectors d (for all n) and which outputs a random variable $\mathcal{G}(d)$ whose support is in $[n]$. We say that \mathcal{S} is a *Shuffler strategy* if it is a function which takes in nonzero n -vectors d (for all n) and outputs a random variable $\mathcal{S}(d)$ whose support is in $\text{supp}(d)$. Intuitively, these definitions say that Guesser is allowed to guess any card type in $[n]$ while Shuffler must select a card type that is still in the deck. For all nonzero d we define the greedy strategy \mathcal{S}_g by

$$\Pr[\mathcal{S}_g(d) = i] = |\text{supp}(d)|^{-1} \quad \forall i \in \text{supp}(d).$$

When n is understood, we let δ_i be the n -vector which has a 1 in position i and 0's everywhere else. We recursively define the score $C(d, \mathcal{G}, \mathcal{S})$ where d is an n -vector and \mathcal{G}/\mathcal{S} are Guesser/Shuffler strategies by having $C(d, \mathcal{G}, \mathcal{S}) = 0$ if $d = (0, \dots, 0)$, and otherwise having

$$C(d, \mathcal{G}, \mathcal{S}) = 1_{\mathcal{S}(d)=\mathcal{G}(d)} + C(d - \delta_{\mathcal{S}(d)}, \mathcal{G}, \mathcal{S}). \quad (2.1)$$

That is, if Shuffler chooses some $i = \mathcal{S}(d)$, then a point is scored if Guesser chooses i as well, and in either case the game continues with the deck $d - \delta_i$ since one copy of i has been removed.

For d an n -vector and \mathcal{S} a Shuffler strategy, define

$$\begin{aligned} f(d, \mathcal{S}) &= \max_{\mathcal{G}} \mathbb{E}[C(d, \mathcal{G}, \mathcal{S})] & f(d) &= \min_{\mathcal{S}} f(d, \mathcal{S}) \\ F(d, \mathcal{S}) &= \min_{\mathcal{G}} \mathbb{E}[C(d, \mathcal{G}, \mathcal{S})] & F(d) &= \max_{\mathcal{S}} F(d, \mathcal{S}), \end{aligned}$$

where the maximum and minimums run through all possible Shuffler and Guesser strategies as appropriate. Intuitively, $f(d, \mathcal{S})$ is the score of the game when Guesser tries to maximize their score given that the deck starts as d and Shuffler uses strategy \mathcal{S} , with the other functions having analogous interpretations. The fact that these maximums and minimums exist can be seen, for instance, by observing that the space of possible strategies forms a compact set. If \mathcal{S} is such that $f(d) = f(d, \mathcal{S})$ for all n -vectors d of any length n , then we say that \mathcal{S} is a *minimizing Shuffler strategy*, and similiary we define what it means for \mathcal{S} to be a *maximizing Shuffler strategy*.

The key lemma that we use throughout this section is the following recurrence relation.

Lemma 2.1. *Let \mathcal{S} be a Shuffler strategy. If d is an n -vector and $p_i = \Pr[\mathcal{S}(d) = i]$, then*

$$\begin{aligned} f(d, \mathcal{S}) &= \max_{i \in [n]} p_i + \sum_{j \in [n]} p_j f(d - \delta_j, \mathcal{S}), \\ F(d, \mathcal{S}) &= \min_{i \in [n]} p_i + \sum_{j \in [n]} p_j f(d - \delta_j, \mathcal{S}). \end{aligned}$$

Proof. Given any Guesser strategy \mathcal{G} , we have by (2.1) that

$$\mathbb{E}[C(d, \mathcal{G}, \mathcal{S})] = \Pr[\mathcal{S}(d) = \mathcal{G}(d)] + \sum p_j \cdot \mathbb{E}[C(d - \delta_j, \mathcal{G}, \mathcal{S})] \leq \max p_i + \sum p_j \cdot f(d - \delta_j, \mathcal{S}),$$

and taking a maximum over \mathcal{G} shows that

$$f(d, \mathcal{S}) = \max_{\mathcal{G}} \mathbb{E}[C(d, \mathcal{G}, \mathcal{S})] \leq \max p_i + \sum p_j \cdot f(d - \delta_j, \mathcal{S}).$$

Moreover, it is straightforward to prove by induction on $\sum d_i$ that the strategy \mathcal{G} which always guesses some k with $p_k = \max p_i$ achieves this upper bound. The proof for $F(d, \mathcal{S})$ is essentially the same and we omit the details. \square

Another key lemma we make use of is the following, which intuitively says that if $\mathcal{S}, \mathcal{S}'$ are two strategies which act the same on d as well as on any deck which is “smaller” than d , then these two strategies give the same expected score if we start with d .

Lemma 2.2. *Let $\mathcal{S}, \mathcal{S}'$ be two shuffler strategies and d an n -vector such that $\mathcal{S}(d') = \mathcal{S}'(d')$ for all n -vectors d' which satisfy $d'_i \leq d_i$ for all i . Then*

$$f(d, \mathcal{S}) = f(d, \mathcal{S}'), \quad F(d, \mathcal{S}) = F(d, \mathcal{S}').$$

Proof. We prove the more general result that, for $d, \mathcal{S}, \mathcal{S}'$ as in the hypothesis, we have

$$\mathbb{E}[C(d, \mathcal{G}, \mathcal{S})] = \mathbb{E}[C(d, \mathcal{G}, \mathcal{S}')] \tag{2.2}$$

for all Guesser strategies \mathcal{G} . With this the stated result following by taking maximum/minimums with respect to this equality.

To prove this general result, assume for the sake of contradiction that there is some d and strategies $\mathcal{S}, \mathcal{S}', \mathcal{G}$ that satisfy the hypothesis of the lemma and such that (2.2) does not hold, and choose such an example with $\sum d_i$ as small as possible. Observe that $d - \delta_i$ for $i \in \text{supp}(d)$ together with $\mathcal{S}, \mathcal{S}'$ also satisfies the hypothesis of the lemma since d does, so by the minimality of d we must have $\mathbb{E}[C(d - \delta_i, \mathcal{G}, \mathcal{S})] = \mathbb{E}[C(d - \delta_i, \mathcal{G}, \mathcal{S}')]$ for all such i . The desired equality now follows from the recursion (2.1). \square

The next result shows the intuitive fact that if, say, $d = (100, 1)$ and Shuffler is trying to minimize their score, then it would be better for them to draw a 1 than a 2 (since the latter will guarantee 100 correct guesses from Guesser).

Lemma 2.3. *Let d be an n -vector. If $d_i \geq d_j > 0$, then*

$$f(d - \delta_i) \leq f(d - \delta_j), \quad F(d - \delta_i) \geq F(d - \delta_j).$$

Proof. Assume for contradiction that there exists such a d with $f(d - \delta_i) > f(d - \delta_j)$, and choose this d so that $\sum d_k$ is as small as possible. Observe that if $d_i = d_j$, then the entries of $d - \delta_i$ and $d - \delta_j$ are permutations of one another, so by the symmetry of the problem we see that $f(d - \delta_i) = f(d - \delta_j)$. Thus we can assume $d_i > d_j$, and in particular $d_i \geq 2$.

Let \mathcal{S} be a minimizing Shuffler strategy and define \mathcal{S}' by having $\mathcal{S}'(d - \delta_i) = \mathcal{S}(d - \delta_j)$ and $\mathcal{S}'(d') = \mathcal{S}(d')$ for all other d' . Observe that this is indeed a Shuffler strategy because

$$\text{supp}(\mathcal{S}'(d - \delta_i)) = \text{supp}(\mathcal{S}(d - \delta_j)) \subseteq \text{supp}(d - \delta_j) \subseteq \text{supp}(d) = \text{supp}(d - \delta_i),$$

where the last equality used $d_i \geq 2$. Let $p_k = \Pr[\mathcal{S}(d - \delta_j) = k]$. We claim that the following is a consequence of Lemmas 2.1 and 2.2:

$$f(d - \delta_i) - f(d - \delta_j) \leq f(d - \delta_i, \mathcal{S}') - f(d - \delta_j, \mathcal{S}) = \sum_{k \in \text{supp}(d - \delta_j)} p_k (f(d - \delta_i - \delta_k) - f(d - \delta_j - \delta_k)). \tag{2.3}$$

In more detail, we have $f(d - \delta_i) \leq f(d - \delta_i, \mathcal{S}')$ because f is a minimum over all possible Shuffler strategies, and $f(d - \delta_j) = f(d - \delta_j, \mathcal{S})$ because \mathcal{S} is a minimizing Shuffler

strategy. This gives the inequality of (2.3). For the equality, we expand $f(d - \delta_i, \mathcal{S}')$ and $f(d - \delta_j, \mathcal{S})$ using Lemma 2.1, noting that $\max_{k \in [n]} \Pr[\mathcal{S}'(d - \delta_i) = k] = \max_{k \in [n]} p_k$ and

$$f(d - \delta_i - \delta_k, \mathcal{S}') = f(d - \delta_i - \delta_k, \mathcal{S}) = f(d - \delta_i - \delta_k),$$

where the first equality uses Lemma 2.2 and the second that \mathcal{S} is a minimizing strategy. This proves (2.3).

Because $d_i > d_j$, we have $(d - \delta_k)_i \geq (d - \delta_k)_j$ for all k , so by the minimality of our choice of d we have $f(d - \delta_k - \delta_i) \leq f(d - \delta_k - \delta_j)$ for all k . Applying this to each term of (2.3) gives $f(d - \delta_i) \leq f(d - \delta_j)$, a contradiction to our assumption on d . This implies the result for f , and an analogous argument gives the result for F . \square

In order for the greedy strategy \mathcal{S}_g to be a minimizing Shuffler strategy, it is necessary that $f(d - \delta_i) \leq 1 + f(d - \delta_j)$ for all d and $i, j \in \text{supp}(d)$. Indeed if this failed for some d , then it would be better for Shuffler to deterministically choose j (giving up a point) than to use any strategy which could draw i . The following shows that a generalization of this condition is true if we inductively assume that the greedy strategy is optimal for all smaller decks.

Lemma 2.4. *Let d be an n -vector such that $f(d', \mathcal{S}_g) = f(d')$ for all $d' \neq d$ which satisfy $d'_k \leq d_k$ for all $k \in [n]$. Then for all $J \subseteq \text{supp}(d)$ and $i \in \text{supp}(d)$, we have*

$$\begin{aligned} |J|f(d - \delta_i) &< 1 + \sum_{j \in J} f(d - \delta_j), \\ |J|F(d - \delta_i) &> 1 + \sum_{j \in J} F(d - \delta_j). \end{aligned}$$

Proof. Let d, i, J be as in the hypothesis of the lemma and assume for contradiction that $|J|f(d - \delta_i) \geq 1 + \sum_{j \in J} f(d - \delta_j)$. Moreover, choose d so that $\sum_k d_k$ is as small as possible, and then choose a set J as small as possible such that this inequality holds for d . By Lemma 2.3, if $d_i \geq d_{j'}$ for some $j' \in J$, then $f(d - \delta_i) \leq f(d - \delta_{j'})$. Thus

$$|J|f(d - \delta_i) \geq 1 + \sum_{j \in J} f(d - \delta_j) \implies |J \setminus \{j'\}|f(d - \delta_i) \geq 1 + \sum_{j \in J \setminus \{j'\}} f(d - \delta_j),$$

contradicting the minimality of J . Thus we must have $d_j > d_i > 0$ for all $j \in J$, and in particular $\text{supp}(d - \delta_j) = \text{supp}(d)$ for all $j \in J$. Let $s = |\text{supp}(d)|$.

First consider the case that $d_i \geq 2$, so $\text{supp}(d - \delta_k) = \text{supp}(d)$ for all $k \in \text{supp}(d)$. Using this and the hypothesis $f(d - \delta_k) = f(d - \delta_k, \mathcal{S}_g)$ for all k , we find by Lemma 2.1 that

$$\begin{aligned} |J|f(d - \delta_i) - \sum_j f(d - \delta_j) &= |J|f(d - \delta_i, \mathcal{S}_g) - \sum_j f(d - \delta_j, \mathcal{S}_g) \\ &= s^{-1} \sum_{k \in \text{supp}(d)} \left(|J|f(d - \delta_i - \delta_k) - \sum_{j \in J} f(d - \delta_j - \delta_k) \right) < s^{-1} \sum_{k \in \text{supp}(d)} 1 = 1, \end{aligned}$$

where the inequality used that each $d - \delta_k$ satisfies the hypothesis of the lemma if d does, and that d is a counterexample to the lemma with $\sum d_p$ as small as possible. This gives the desired inequality.

Thus we can assume $d_i = 1$. Consider the strategy \mathcal{S}' which agrees with \mathcal{S}_g on $d' \neq d - \delta_i$ and which has

$$\Pr[\mathcal{S}'(d - \delta_i) = j] = \begin{cases} s^{-1}(|J|^{-1} + 1) & j \in J, \\ s^{-1} & j \in \text{supp}(d) \setminus (J \cup \{i\}). \end{cases}$$

Note that these probabilities sum to 1 and that this is indeed a Shuffler strategy. By Lemmas 2.1 and 2.2 we have

$$\begin{aligned} |J|f(d - \delta_i) &\leq |J|f(d - \delta_i, \mathcal{S}') \\ &= s^{-1}(1 + |J|) + s^{-1} \sum_{j \in J} (1 + |J|)f(d - \delta_i - \delta_j) \\ &\quad + s^{-1} \sum_{k \in \text{supp}(d) \setminus (J \cup \{i\})} |J|f(d - \delta_i - \delta_k) \\ &= s^{-1} + s^{-1}|J| + s^{-1} \sum_{j \in J} f(d - \delta_i - \delta_j) + s^{-1} \sum_{k \in \text{supp}(d) \setminus \{i\}} |J|f(d - \delta_i - \delta_k). \end{aligned}$$

On the other hand, we have

$$\begin{aligned} \sum_{j \in J} f(d - \delta_j) &= s^{-1}|J| + s^{-1} \sum_{j \in J} \left(f(d - \delta_j - \delta_i) + \sum_{k \in \text{supp}(d) \setminus \{i\}} f(d - \delta_j - \delta_k) \right) \\ &= s^{-1}|J| + s^{-1} \sum_{j \in J} f(d - \delta_j - \delta_i) + s^{-1} \sum_{k \in \text{supp}(d) \setminus \{i\}} \sum_{j \in J} f(d - \delta_j - \delta_k) \end{aligned}$$

and subtracting these two expressions gives

$$\begin{aligned} |J|f(d - \delta_i) - \sum_{j \in J} f(d - \delta_j) &\leq s^{-1} + s^{-1} \sum_{k \in \text{supp}(d) \setminus \{i\}} \left(|J|f(d - \delta_i - \delta_k) - \sum_{j \in J} f(d - \delta_j - \delta_k) \right) \\ &< s^{-1} + s^{-1} \cdot (s - 1) = 1, \end{aligned}$$

where again this inequality used the minimality of d . This completes the proof for f , and a nearly identical argument works for F . \square

With this we can prove our main result for this section, which is a slightly more precise version of Theorem 1.3. To state our result, we say that an n -vector d is *fully-supported* if $\text{supp}(d) = [n]$.

Theorem 2.5. *Let \mathcal{S} be a Shuffler strategy and \mathcal{S}_g the greedy strategy. Then \mathcal{S} is a minimizing Shuffler strategy if and only if $\mathcal{S}(d)$ has the same distribution as $\mathcal{S}_g(d)$ for all n -vectors d , and \mathcal{S} is a maximizing Shuffler strategy if and only if $\mathcal{S}(d)$ has the same distribution as $\mathcal{S}_g(d)$ for all fully-supported n -vectors d .*

Proof. Let \mathcal{S} be a minimizing Shuffler strategy, and assume for contradiction that there exists an n -vector d such that $\mathcal{S}(d)$ does not have the same distribution as $\mathcal{S}_g(d)$. Choose such a d so that $\sum d_i$ is as small as possible, which implies that Lemma 2.4 applies to d .

For all k , let $p_k = \Pr[\mathcal{S}(d) = k]$. By our choice of d , we have $p_k \neq p_{k'}$ for some $k, k' \in \text{supp}(d)$. By relabeling the card types, we can assume

$$p_1 = p_2 = \dots = p_i > p_{i+1} \geq p_{i+2} \geq \dots \geq p_n$$

for some $i \geq 1$ with $i + 1 \in \text{supp}(d)$. Define $\varepsilon = (p_1 - p_{i+1})/(1 + 1/i)$ and let \mathcal{S}' be the strategy which has $\mathcal{S}'(d') = \mathcal{S}(d')$ for all $d' \neq d$, and for d we have

$$\Pr[\mathcal{S}'(d) = k] = \begin{cases} p_1 - i^{-1}\varepsilon & k \leq i, \\ p_{i+1} + \varepsilon & k = i + 1, \\ p_k & k > i + 1. \end{cases}$$

Note that this is a Shuffler strategy because these probabilities add to 1 and $\Pr[S'(d) = k] > 0$ only if $k \in \text{supp}(d)$. Also with this we have

$$\max_k \Pr[S'(d) = k] = p_{i+1} + \varepsilon = p_1 - i^{-1}\varepsilon.$$

Using Lemmas 2.1 and 2.2 and that \mathcal{S} is a minimizing Shuffler strategy, we find

$$f(d, \mathcal{S}') = p_1 - i^{-1}\varepsilon + \sum_{k \leq i} (p_k - i^{-1}\varepsilon) f(d - \delta_k) + (p_{i+1} + \varepsilon) f(d - \delta_{i+1}) + \sum_{k > i+1} p_k f(d - \delta_k).$$

Combining this with the recurrence for \mathcal{S} gives

$$f(d, \mathcal{S}') - f(d, \mathcal{S}) = -i^{-1}\varepsilon - i^{-1}\varepsilon \sum_{k \leq i} f(d - \delta_k) + \varepsilon f(d - \delta_{i+1}) < 0,$$

where this last inequality used Lemma 2.4 with $i+1$ and $J = \{1, 2, \dots, i\}$. This contradicts \mathcal{S} being a minimizing Shuffler strategy, giving the desired result.

The proof for the maximizing case is almost identical, and we briefly sketch the ideas for this argument. If d is not fully-supported then $F(d, \mathcal{S}) = 0$ for all Shuffler strategies, so \mathcal{S} is a maximizing Shuffler strategy if and only if $F(d, \mathcal{S}) = F(d)$ for all fully-supported d . We assume for contradiction that \mathcal{S} is a maximizing Shuffler strategy and that there exists a fully-supported d such that $\mathcal{S}(d)$ does not have the same distribution as $\mathcal{S}_g(d)$. This means there exists an i with

$$p_1 = p_2 = \dots = p_i < p_{i+1} \leq p_{i+2} \leq \dots \leq p_n.$$

We then define the strategy \mathcal{S}' exactly as before except that we use $\varepsilon = -(p_1 - p_{i+1})/(1 + 1/i)$. Here d being fully-supported is essential, as otherwise we might have $1 \notin \text{supp}(d)$ and that $\Pr[S'(d) = 1] = -i^{-1}\varepsilon > 0$, so this would not be a Shuffler strategy. From here the same analysis as before gives the result. \square

3 Proof of Theorem 1.4

Given the work of the previous section, to prove Theorem 1.4, it suffices to determine how well the greedy strategy does against an optimal Guesser strategy. A key fact is that every “reasonable” Guesser strategy performs equally well against the greedy strategy in expectation.

Lemma 3.1. *If \mathcal{G} is a Guesser strategy which always guesses a card type that is in the deck, then*

$$\mathbb{E}[C_{m,n}(\mathcal{G}, \mathcal{S}_g)] = \max_{\mathcal{G}'} \mathbb{E}[C_{m,n}(\mathcal{G}', \mathcal{S}_g)].$$

If \mathcal{G} is a Guesser strategy which, whenever possible, guesses a card type which is not in the deck, then

$$\mathbb{E}[C_{m,n}(\mathcal{G}, \mathcal{S}_g)] = \min_{\mathcal{G}'} \mathbb{E}[C_{m,n}(\mathcal{G}', \mathcal{S}_g)].$$

This result is straightforward to prove by using Lemma 2.1 and induction on $\sum d_i$, so we omit its proof. With this lemma we can prove the bounds of Theorem 1.4 by analyzing the performance of any specific Guesser strategy which is “reasonable”. In particular, we define \mathcal{G}^+ to be some Guesser strategy which always guesses a card type which appears with the highest multiplicity in the deck.

Let C_k denote the number of times Guesser correctly guesses a card type which had multiplicity k in the deck. Note that this is a random variable which depends on the choices of strategies \mathcal{G}, \mathcal{S} in the game, and also note that we always have $C_{m,n}(\mathcal{G}, \mathcal{S}) = \sum C_k$. Let t_k be the smallest integer such that at the start of round t_k , each card type

has fewer than k copies left in the deck. Define $V_{k,\ell}$ to be the number of card types which have at least ℓ copies left in the deck at the start of round t_k . For example, $V_{k,\ell} = 0$ for all $k \leq \ell$ and $V_{k,\ell} \leq V_{k',\ell}$ for all $k \leq k'$. We also adopt the convention that $V_{m+1,m} = n$.

Lemma 3.2. *If Guesser and Shuffler use the strategies \mathcal{G}^+, S_g in the Guesser-Shuffler game with a deck which has n card types of multiplicity m , then*

$$\mathbb{E}[C_m] = \log n + O(1),$$

and for $k < m$ we have

$$\mathbb{E}[C_k] \leq \log(\mathbb{E}[V_{k+1,k}]) + O(1)$$

.

Note that $V_{k+1,k} \geq 1$, so this logarithm is well defined.

Proof. We claim that for all k ,

$$\mathbb{E}[C_k | V_{k+1,k} = r] = \sum_{x=1}^r x^{-1}.$$

Indeed, after the last card type of multiplicity $k+1$ is drawn, Guesser (using strategy \mathcal{G}^+) will continue to guess the remaining r card types of multiplicity k until they are depleted. Because Shuffler uses the greedy strategy, the probability that the first of these card types drawn matches what Guesser guesses is exactly r^{-1} , the probability the second matches is $(r-1)^{-1}$ and so on; so the claim follows by linearity of expectation. Thus

$$\mathbb{E}[C_k] = \sum_{r \geq 1} \Pr[V_{k+1,k} = r] \cdot \sum_{x=1}^r x^{-1} = \sum_{r \geq 1} \Pr[V_{k+1,k} = r] \cdot (\log r + O(1)) = \mathbb{E}[\log V_{k+1,k}] + O(1).$$

The $k = m$ case of the lemma follows because $V_{m+1,m} = n$ deterministically, and the $k < m$ result follows from Jensen's inequality since \log is a convex function. \square

We next show the following.

Lemma 3.3. *If Shuffler uses the greedy strategy S_g in the Guesser-Shuffler game with a deck which has n card types of multiplicity m , then for all $k < m$ we have*

$$\mathbb{E}[V_{k+1,k}] \leq \mathbb{E}[V_{m,k}] \sim \frac{(\log n)^k}{k!}$$

Proof. The first inequality follows from $V_{k+1,k} \leq V_{m,k}$ since $k+1 \leq m$, so it suffices to prove the asymptotic result. We do this by showing that $V_{m,k}$ has the same distribution as a random variable related to a variant of the coupon collector problem.

The variant goes as follows: every day a household with m brothers b_1, \dots, b_m receives a coupon c . There are n different types of coupons, each appearing with probability $1/n$. If b_m does not already own a copy of c , then he adds it to his collection. Otherwise he gives it to his younger brother b_{m-1} who keeps c if he needs it, and otherwise he gives it to b_{m-2} and so on. Observe that b_m will always be the first to obtain every type of coupon. Let U_k denote the number of coupons that b_k is missing when b_m completes their collection.

We claim that U_k has the same distribution as $V_{m,k}$. Indeed, we can form a coupling as follows. Let c_1, c_2, \dots be the sequence of coupons that appear, and let $f(t)$ denote the t th coupon that was used in one of the m collections (that is, this is the t th coupon after we throw away any coupon which has already appeared more than m times). We then define a deck of cards π by letting $\pi_t = c_{f(t)}$ for all t . Because each coupon type is equally likely to appear each day, it is not hard to show that π is distributed the same as

if Shuffler made it using the greedy strategy \mathcal{S}_g . Moreover, having b_k missing r coupon types when b_m finishes their set is equivalent to π having r card types with multiplicity k after the last card with multiplicity m is drawn in π , proving the claim.

It was proven by Foata, Han, and Lass [11, Equation (10.1)] that

$$\mathbb{E}[U_k] \sim \frac{(\log n)^k}{k!}$$

for any fixed k , proving the result. See also Adler, Oren and Ross [1] for a discussion of this result in English. This completes the proof. \square

With this we can now prove our main result.

Proof of Theorem 1.4. By Theorem 1.3 and Lemma 3.1, we see that proving (1.1) is equivalent to showing that

$$\mathbb{E}[C_{m,n}(\mathcal{G}^+, \mathcal{S}_g)] \sim \log n,$$

where \mathcal{G}^+ is the strategy of guessing a card type which has highest multiplicity in the deck. This asymptotic result then follows by Lemmas 3.2 and 3.3 together with the trivial bound $C_k \geq 0$ for all k .

To prove (1.2), let \mathcal{G} be any Guesser strategy which stops guessing card types that are in the deck once a card type runs out in the deck. by Theorem 1.3 and Lemma 3.1, it suffices to prove

$$\mathbb{E}[C_{m,n}(\mathcal{G}, \mathcal{S}_g)] \sim \Gamma\left(1 + \frac{1}{m}\right) (m!)^{1/m} \cdot n^{-1/m}.$$

Let π be the deck shuffled according to \mathcal{S}_g , and let I_t be the indicator function which is 1 if \mathcal{G} correctly guesses π_t . Let T denote the largest index t such that $\{\pi_1, \dots, \pi_{t-1}\}$ does not contain a card type with multiplicity m . Then $I_t = 0$ if $t > T$, and more generally we have

$$\Pr[I_t = 1] = n^{-1} \cdot \Pr[T \geq t],$$

as each card type is equally likely to appear in round $t \leq T$. Because $C_{m,n}(\mathcal{G}, \mathcal{S}_g) = \sum I_t$, we see that

$$\mathbb{E}[C_{m,n}(\mathcal{G}, \mathcal{S}_g)] = n^{-1} \sum \Pr[T \geq t] = n^{-1} \cdot \mathbb{E}[T], \tag{3.1}$$

so it suffices to compute this expectation.

We can interpret T in terms of a variant of the birthday problem. Namely, we can consider the following experiment: sample people with replacement until we find m people which all have the same birthday. If there are n possible birthdays and each are equally likely, then the number of people we sample has the same distribution as T (which can be proved using a coupling argument as in Lemma 3.3). The expected value of this statistic for the birthday problem was computed by Klamkin and Newman [16, Theorem 2], and their result implies

$$\mathbb{E}[T] = \Gamma\left(1 + \frac{1}{m}\right) (m!)^{1/m} \cdot n^{1-1/m},$$

and plugging this into (3.1) gives the desired result. \square

4 Proof of Proposition 1.5

We formally define the Restricted Matching Pennies game as follows. We define n -vectors and δ_i as in Section 2. We say that \mathcal{A} is a strategy in this game if it is a function which takes in two n -vectors a, b (in this order) and outputs a random variable $\mathcal{A}(a, b)$ whose support lies in $\text{supp}(a)$. We recursively define the score $M(a, b, \mathcal{A}, \mathcal{B})$ where a, b

are n -vectors with $\sum a_i = \sum b_i$ and \mathcal{A}, \mathcal{B} are strategies by being 0 if $a = b = (0, \dots, 0)$, and otherwise having

$$M(a, b, \mathcal{A}, \mathcal{B}) = 1_{\mathcal{A}(a,b)=\mathcal{B}(b,a)} + M(a - \delta_{\mathcal{A}(a,b)}, b - \delta_{\mathcal{B}(b,a)}, \mathcal{A}, \mathcal{B}).$$

Our main goal for this section is to prove the following strengthening of Proposition 1.5 which works for decks with arbitrary multiplicities.

Proposition 4.1. *Let \mathcal{U} be the strategy with $\Pr[\mathcal{U}(a, b) = i] = a_i / \sum a_j$ for all n -vectors a, b . Then for any strategies \mathcal{A}, \mathcal{B} and any two n -vectors a, b with $\sum a_i = \sum b_i = N$, we have*

$$\mathbb{E}[M(a, b, \mathcal{A}, \mathcal{U})] = \mathbb{E}[M(a, b, \mathcal{U}, \mathcal{B})] = \sum_{i \in [n]} \frac{a_i b_i}{N}.$$

Proof. For some slight ease of notation we define $f(a, b, \mathcal{A}) = \mathbb{E}[M(a, b, \mathcal{A}, \mathcal{U})]$. We prove by induction on N that for all n -vectors a, b with $\sum a_i = \sum b_i = N$, we have

$$f(a, b, \mathcal{A}) = \sum_{i \in [n]} \frac{a_i b_i}{N}$$

for all strategies \mathcal{A} , with the analogous result for \mathcal{B} being proven in exactly the same way.

The base case $N = 0$ is trivial. Assume we have proven the result up to some value N and let a, b be n -vectors with $\sum a_i = \sum b_i = N$. Let

$$f_{i,j}(a, b, \mathcal{A}) = \mathbb{E}[M(a, b, \mathcal{A}, \mathcal{U}) | \mathcal{A}(a, b) = i, \mathcal{U}(b, a) = j]$$

and define $p_i = \Pr[\mathcal{A}(a, b) = i]$. Because \mathcal{U} guesses j with probability b_j/N , we see that

$$f(a, b, \mathcal{A}) = \sum_i p_i \sum_j b_j N^{-1} \cdot f_{i,j}(a, b, \mathcal{A}). \tag{4.1}$$

By our inductive hypothesis and the definition of M , we have

$$f_{i,i}(a, b, \mathcal{A}) = 1 + \frac{(a_i - 1)(b_i - 1)}{N - 1} + \sum_{k \neq i} \frac{a_k b_k}{N - 1} = (N - 1)^{-1} (N - b_i - a_i + \sum_k a_k b_k),$$

and for $i \neq j$ we have

$$f_{i,j}(a, b, \mathcal{A}) = \frac{(a_i - 1)b_i}{N - 1} + \frac{a_j(b_j - 1)}{N - 1} + \sum_{k \neq i,j} \frac{a_k b_k}{N - 1} = (N - 1)^{-1} (-b_i - a_j + \sum_k a_k b_k).$$

Thus for any fixed i we have

$$\begin{aligned} \sum_j b_j (N - 1) \cdot f_{i,j}(a, b, \mathcal{A}) &= N b_i + \sum_j b_j \left(-b_i - a_j + \sum_k a_k b_k \right) \\ &= N b_i - N b_i - \sum_j a_j b_j + N \sum_k a_k b_k = (N - 1) \sum_j a_j b_j, \end{aligned}$$

where the second equality used that $\sum b_j = N$ in two places. Plugging this expression into (4.1) and using $\sum_i p_i = 1$ completes the inductive step in the proof of $f(a, b, \mathcal{A}) = \sum a_j b_j / N$ for all strategies \mathcal{A} , proving the result. \square

5 Concluding remarks

Strategies for Guesser. In Theorem 1.3, we showed that the greedy strategy is optimal for Shuffler in the Guesser-Shuffler game. It is natural to ask what the optimal strategies are for Guesser. Computations suggest that such strategies are complicated to describe, even for $n = 2$ or $m = 2$. However, we do know some simple strategies for Guesser which perform close to best possible. For example, in the introduction we noted that the Guesser strategy \mathcal{G}^+ which chooses a card type with highest multiplicity each round asymptotically gives $\log n$ correct guesses in expectation, which is best possible.

Similarly one can come up with a strategy \mathcal{G}^- which achieves at most $\Theta_m(n^{-1/m})$ correct guesses in expectation. This strategy goes through several phases, starting in Phase m . If the strategy is currently in Phase i with $i > 0$ and there are at least $n^{(i-1)/m}$ card types in the deck with multiplicity smaller than i , then proceed to Phase $i - 1$. Otherwise uniformly at random guess a card type which appears with multiplicity i . When the game reaches Phase 0 there is some card type which is no longer in the deck, and we guess this card type for the rest of the game. It is not difficult to show that the expected score during each Phase is at most $n^{-1/m}$, so for m fixed this gives about $n^{-1/m}$ correct guesses.

Partial Feedback. The Guesser-Shuffler game is an adversarial version of the complete feedback model, and there are natural adversarial versions of other variants of this model. For example, in the partial feedback model, the Guesser is only told whether their guess was correct or not each round (cf. this with the complete feedback model where the card is always revealed to the Guesser each round). Motivated by this, we define the *online partial feedback game* to be the two player game where each round Shuffler chooses a card from the deck, Guesser guesses it, and then Guesser is told whether their guess is correct or not. We similarly define the *offline partial feedback game* where now Shuffler can freely shuffle the deck at the start of the game, but afterwards can not change the order of the cards.

We note that in the Guesser-Shuffler game, the distinction between online and offline games were effectively irrelevant. This is roughly because the state of the game (but not the score) at any point in time depends only on Shuffler's choices (i.e. the cards which they discarded from the deck), so Guesser's actions during the game should not influence how Shuffler shuffles. However, under partial feedback, the state of the game also depends on Guesser's actions (i.e. this determines how much information Guesser has at any point), and seemingly Shuffler should be able to utilize this information during the game.

It was shown in [7] that for n significantly large in terms of m , Guesser can not get more than $m + O(m^{3/4} \log m)$ points in expectation when the deck is shuffled uniformly at random and the Guesser is given partial feedback. Thus in the adversarial models one can not hope for Guesser to get asymptotically more than m correct guesses, which they can always guarantee by guessing the same card type each round. When the deck is shuffled uniformly at random, it is known [8] that Guesser can play so that they get $m + \Omega(m^{1/2})$ correct guesses when n is sufficiently large, but we do not know of any such strategy that works in the adversarial setting.

Question 5.1. Is there a strategy for Guesser in either the online or offline partial feedback game which uses a deck with m copies of n different card types such that their expected score is "significantly" larger than m ?

The best strategy we are aware of in the offline version is for Guesser to uniformly at random choose some $i_1 \in [n]$, then guess i_1 each round until they get m correct guesses, then to uniformly at random choose some $i_2 \in [n] \setminus \{i_1\}$ and guess i_2 until m correct guesses are made, and so on. It is not difficult to see that Shuffler can arrange the

deck ahead of time so that Guesser gets $m + 1/2$ correct guesses in expectation under this strategy when $m > 1$ (namely by shuffling the deck so that the last n cards are all distinct). However, in the online version of this game, Shuffler can guarantee that Guesser gets exactly m correct guesses using this strategy when $m > 1$.

The best strategy we are aware of in the online setting is for Guesser to uniformly at random guess either 1 or 2 each round until they correctly guess either m 1's or m 2's, at which point they only guess the other card type for the rest of the game. Note that if Guesser simply uniformly guessed either 1 or 2 each round, then it is easy to show that they get m points in expectation. By additionally ignoring cards that Guesser knows are not in the deck, one can show that they improve their score by roughly 2^{-m} in expectation.

Other Games. One can form a common generalization of the Guesser-Shuffler and Restricted Matching Pennies games as follows. For any integer k , define the k -Guesser-Shuffler game to be the same as the Guesser-Shuffler game with the restriction that Guesser can guess each card type at most k times throughout the game. For example, if each of the n card types appears m times, then the mn -Guesser-Shuffler game is equivalent to the Guesser-Shuffler game (since no restriction is placed on Guesser), and the m -Guesser-Shuffler game is equivalent to the Restricted Matching Pennies game.

Question 5.2. What are the optimal strategies in the k -Guesser-Shuffler game when one player wishes to maximize and the other wishes to minimize the number of correct guesses made during this game? How many correct guesses are made in expectation under these strategies?

Our results in Theorems 1.3, 1.4, and Proposition 1.5 shows that the answers to this question vary significantly between $k = mn$ and $k = m$, and it would be interesting to know if there is some sort of interpolation between these results as k varies.

The k -Guesser-Shuffler game can be viewed as playing multiple rounds of the classical Matching Pennies game where restrictions are placed on how many times each player may use each move. It is natural to consider what happens for other games. For example, the game Rock, Paper, Scissors has two players selecting one of Rock, Paper, or Scissors each round with Rock beating Scissors, Scissors beating Paper, and Paper beating Rock. If a player Alice picks a move that beats the move of Bob, then Alice is given a point and Bob loses a point, and if they pick the same move both players are given 0 points.

Question 5.3. Define the game Semi-Restricted Rock, Paper, Scissors by playing $3m$ rounds of Rock, Paper, Scissors where one player must use each move exactly m times (with no restrictions placed on the other player). What are the optimal strategies in this game and what is the expected score under these strategies?

We note that similar games have appeared in the manga *Tobaku Mokushiroku Kaiji* [12], namely the games “Restricted Rock, Paper, Scissors” and “E-Card.”

References

- [1] I. Adler, S. Oren, and S. M. Ross. The coupon-collector's problem revisited. *Journal of Applied Probability*, pages 513–518, 2003. MR1978107
- [2] D. Aldous and P. Diaconis. Shuffling cards and stopping times. *The American Mathematical Monthly*, 93(5):333–348, 1986. MR0841111
- [3] D. Bayer and P. Diaconis. Trailing the dovetail shuffle to its lair. *The Annals of Applied Probability*, pages 294–313, 1992. MR1161056
- [4] D. Blackwell and J. Hodges Jr. Design for the control of selection bias. *The Annals of Mathematical Statistics*, pages 449–460, 1957. MR1105868
- [5] M. Ciucu. No-feedback card guessing for dovetail shuffles. *Annals of Applied Probability*, 8(4):1251–1269, 1998. MR1661184

- [6] P. Diaconis and R. Graham. The analysis of sequential experiments with feedback to subjects. *The Annals of Statistics*, 9(1):3–23, 1981. MR0600529
- [7] P. Diaconis, R. Graham, X. He, and S. Spiro. Card guessing with partial feedback. *Combinatorics, Probability and Computing*, page 1–20, 2021. MR4356453
- [8] P. Diaconis, R. Graham, and S. Spiro. Guessing about guessing: Practical strategies for card guessing with feedback. *arXiv preprint arXiv:2012.04019*, 2020.
- [9] P. Diaconis and M. Shahshahani. Generating a random permutation with random transpositions. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 57(2):159–179, 1981.
- [10] B. Efron. Forcing a sequential experiment to be balanced. *Biometrika*, 58(3):403–417, 1971. MR0312660
- [11] D. Foata and G.-N. Han. Les nombres hyperharmoniques et la fratrie du collectionneur de vignettes. *Seminaire Lotharingien de Combinatoire*, 47:20pp, 2001. MR1894021
- [12] N. Fukumoto. *Tobaku Mokushiroku Kaiji*. Weekly Young Magazine, 1996.
- [13] E. Gilbert. Theory of shuffling. Technical report, Technical memorandum, Bell Laboratories, 1955.
- [14] J. He and A. Ottolini. Card guessing and the birthday problem for sampling without replacement. *arXiv:2108.07355*
- [15] J. Jonasson. The overhand shuffle mixes in $\theta(n^2 \log n)$ steps. *The Annals of Applied Probability*, 16(1):231–243, 2006. MR2209341
- [16] M. S. Klamkin and D. J. Newman. Extensions of the birthday surprise. *Journal of Combinatorial Theory*, 3(3):279–282, 1967. MR0224121
- [17] T. Krityakierne and T. A. Thanatipanonda. The card guessing game: A generating function approach. *arXiv:2107.11142*
- [18] P. Liu. On card guessing game with one time riffle shuffle and complete feedback. *Discrete Applied Mathematics*, 288:270–278, 2021. MR4156915
- [19] L. Pehlivan. No feedback card guessing for top to random shuffles. *arXiv:1006.1321*

Acknowledgments. The author would like to thank Jimmy He for pointing out minor typos in an earlier draft, as well as the anonymous referee who offered a number of useful suggestions.

Electronic Journal of Probability

Electronic Communications in Probability

Advantages of publishing in EJP-ECP

- Very high standards
- Free for authors, free for readers
- Quick publication (no backlog)
- Secure publication (LOCKSS¹)
- Easy interface (EJMS²)

Economical model of EJP-ECP

- Non profit, sponsored by IMS³, BS⁴, ProjectEuclid⁵
- Purely electronic

Help keep the journal free and vigorous

- Donate to the IMS open access fund⁶ (click here to donate!)
- Submit your best articles to EJP-ECP
- Choose EJP-ECP over for-profit journals

¹LOCKSS: Lots of Copies Keep Stuff Safe <http://www.lockss.org/>

²EJMS: Electronic Journal Management System <http://www.vtex.lt/en/ejms.html>

³IMS: Institute of Mathematical Statistics <http://www.imstat.org/>

⁴BS: Bernoulli Society <http://www.bernoulli-society.org/>

⁵Project Euclid: <https://projecteuclid.org/>

⁶IMS Open Access Fund: <http://www.imstat.org/publications/open.htm>