Semi-restricted Rock, Paper, Scissors

Sam Spiro[a] Erlang Surya[a] Ji Zeng[a]

Submitted: Sep 23, 2022; Accepted: Oct 5, 2023; Published: Dec 1, 2023

Abstract

Consider the following variant of Rock, Paper, Scissors (RPS) played by two players Rei and Norman. The game consists of $3n$ rounds of RPS, with the twist being that Rei (the restricted player) must use each of Rock, Paper, and Scissors exactly n times during the $3n$ rounds, while Norman is allowed to play normally without any restrictions. Answering a question of Spiro, we show that a certain greedy strategy is the unique optimal strategy for Rei in this game, and that Norman's expected score is $\Theta(\sqrt{n})$. Moreover, we study semi-restricted versions of general zero sum games and prove a number of results concerning their optimal strategies and expected scores, which in particular implies our results for semi-restricted RPS.

Mathematics Subject Classifications: 91A05

1 Introduction

1.1 Rock, Paper, Scissors

The game Rock, Paper, Scissors, or RPS for short, is a popular game whose first known usage dates back to China nearly 2000 years ago. A round of RPS consists of two players simultaneously selecting to play either Rock, Paper, or Scissors; where Rock beats Scissors, Scissors beats Paper, and Paper beats Rock. If a player uses a move which beats their opponents move, then that player gains a point and their opponent loses a point, with the match resulting in a draw if both players select the same move.

RPS, while fun to play, is not particularly interesting from a mathematical perspective. As a symmetric zero-sum game, both players have an expected score of 0, and it is easy to show that the unique optimal strategy for both players is to choose each of the three options with uniform probability; see for example [12, Chapter 17]. In this paper we study a non-trivial variant of RPS, inspired by similar games introduced by Fukumoto [4] and Spiro [11], called *semi-restricted RPS*.

Semi-restricted RPS is a perfect-information zero-sum game played by two players named Rei and Norman. The game consists of $3n$ rounds of RPS, with the twist being

[a]Department of Mathematics, University of California San Diego, La Jolla, California, U.S.A. (sspiro@ucsd.edu, esurya@ucsd.edu, jzeng@ucsd.edu).

that Rei must use each of Rock, Paper, and Scissors exactly n times during the $3n$ rounds, while Norman is allowed to play normally without any restrictions. It is clear that Norman has an advantage in this game, since in particular he is guaranteed to win the last round of the game (assuming he is paying attention and has a good memory). However, it is unclear how much more Norman is expected to win over Rei when both play optimally, and it is also unclear what the optimal strategies are for either player.

One simple strategy that Rei can implement is what we call the *greedy strategy*. Under this strategy, if Rei can still play each of Rock, Paper, and Scissors, she selects each option with probability 1/3, regardless of how many actions remain of each option. If she can only play, say, Rock and Paper, then she chooses Paper with probability 2/3 and Rock with probability 1/3. More generally, if she has two options remaining, she will choose the option which beats the other with probability 2/3 and the other option with probability 1/3. And of course, if only one option remains, she plays this with probability 1.

This strategy is "greedy" because with this strategy, Rei minimizes her expected loss for any given round. Indeed, if she has all three options remaining, then playing each with probability 1/3 makes it so that, regardless of what Norman does, Rei is just as likely to win as she is to lose. Similarly if only Rock and Paper remains, then Norman should only ever play Paper or Scissors (since Rock is guaranteed not to win). If Rei follows the greedy strategy and Norman plays Paper, then Rei will lose with probability 1/3. If Norman plays Scissors, then Rei will lose with probability 2/3 but win with probability 1/3. Thus regardless of what Norman does, Rei will expect to lose 1/3 points each round under the greedy strategy, and one can show that this is best possible.

While the greedy strategy is optimal if Rei is only concerned about a given round, it is far from clear that this is a good strategy overall. Indeed, say the game reaches the point where Rei can play 100 Rocks, 100 Papers, and just 1 Scissors. Intuitively, in this scenario Rei should not play Scissors with high probability, as doing so will severely limit her remaining options for the rest of the game. As such the following may come as a bit of a surprise.

Theorem 1. *The greedy strategy is the unique optimal strategy for Rei in semi-restricted RPS. Moreover, if the game consists of* $3n$ *rounds with both players playing optimally, then Norman's expected score is* $\Theta(\sqrt{n})$.

Here and throughout the paper we make use of standard asymptotic notation, see Subsection 1.3 for precise definitions.

1.2 A More General Setting

The problem of studying semi-restricted RPS was first proposed by Spiro [11]. In [11], a semi-restricted version of another classical zero-sum game, Matching Pennies, was introduced. This semi-restricted game was motivated by a certain card guessing game studied by Diaconis and Graham [2] which has recently received a fair amount of attention, see for example [1, 3, 6, 7, 9].

While it is easy to generalize these examples to study "semi-restricted versions" of arbitrary simultaneous zero-sum games, for simplicity we will focus on games which come

from digraphs. However, most of our results hold in broader generality. Recall that a *directed graph* (*digraph* shortly) D consists of a set of vertices $V(D)$ together with a set of ordered pairs of vertices $E(D)$ called *arcs*. We will sometimes denote arcs (u, v) as $u \to v$ or uv. Throughout this paper we will only consider digraphs without self-loops and which have at most one arc between a given pair of vertices.

Given a digraph D, we define the *D-game* by having two players simultaneously select a vertex of D. If the selected vertices are u, v and $uv \in E(D)$, then the player who chose u gains a point and the player who chose v loses a point, and nothing happens if $uv, vu \notin E(D)$. For example, if D is the circuit of length 3, i.e. the 3-vertex digraph with arcs $1 \to 2 \to 3 \to 1$, then the D-game is equivalent to RPS. Observe that the D-game is always a symmetric zero-sum game.

We will say that a vector r is a *restriction vector* (with respect to a digraph D) if it is a vector of non-negative integers indexed by $V(D)$. Given a digraph D and restriction vector r, we define the *semi-restricted D-game with parameter r* by having two players Rei and Norman iteratively play the D-game for a total of $\sum_u r_u$ rounds with the restriction that Rei must select each $u \in V(D)$ exactly r_u times. We let[1] $S_D(r)$ denote the expected score for Norman when both players play optimally in the semi-restricted D game with parameter r, and throughout we let $\mathbf{1}$ denote the all 1's vector of dimension $|V(D)|$. For example, Theorem 1 says $S_D(n \cdot \mathbf{1}) = \Theta(\sqrt{n})$ when D is the circuit of length 3.

Determining optimal strategies for semi-restricted D-games in full generality seems impossible. Nevertheless, we are able to obtain effective bounds on $S_D(r)$ for all D. To state these results, given a digraph D and $v \in V(D)$, we define the *out-neighborhood* $N^+(v) = \{u : (v, u) \in E(D)\}$, and similarly we define the *in-neighborhood* $N^-(v) = \{u : (u, v) \in E(D)\}$. We let $d^+(v) = |N^+(v)|$ and $d^-(v) = |N^-(v)|$ denote the *out-degree* and *in-degree* of v, respectively.

A basic observation is that for all vertices $v \in V(D)$, we have

$$S_D(n \cdot \mathbf{1}) \geqslant (d^+(v) - d^-(v))n.$$

Indeed, if Norman uses the deterministic strategy of playing v every round, then he will win exactly $d^+(v)n$ rounds and lose exactly $d^-(v)n$ rounds, so he can achieve an expected score of at least $(d^+(v) - d^-(v))n$ with this strategy. It turns out that this trivial lower bound is close to best possible.

Theorem 2. *For all digraphs D and $n \geqslant 1$, we have*

$$\max_v \{d^+(v) - d^-(v)\}n \leqslant S_D(n \cdot \mathbf{1}) \leqslant \max_v \{d^+(v) - d^-(v)\}n + O_D(\sqrt{n}).$$

Similar bounds hold for general restriction vectors r, though in this case our bound on the error term is weaker.

Theorem 3. *For all digraphs D and restriction vectors r, we have*

$$\max_v \left\{ \sum_{u \in N^+(v)} r_u - \sum_{u \in N^-(v)} r_u \right\} \leqslant S_D(r) \leqslant \max_v \left\{ \sum_{u \in N^+(v)} r_u - \sum_{u \in N^-(v)} r_u \right\} + O_D(M^{2/3}),$$

[1] We will define these terms more formally in Section 2.

where $M = \max_v r_v$.

The bounds of Theorem 2 are asymptotically tight except when $d^+(v) = d^-(v)$ for all v, and when this happens we say that D is *Eulerian*. A class of Eulerian digraphs that are of particular interest to us are Eulerian tournaments, where a *tournament* is a digraph such that either $uv \in E(D)$ or $vu \in E(D)$ is in D for all distinct $u, v \in E(D)$ (equivalently, tournaments are orientations of complete graphs). For example, the circuit of length 3 is an Eulerian tournament. More generally, any reasonable extension of RPS (such as "Rock, Paper, Scissors, Lizard, Spock" or the infamous "RPS-25" consisting of 25 different options) will have a corresponding digraph which is an Eulerian tournament.

By Theorem 1, we know that semi-restricted Rock, Paper, Scissors has a simple greedy strategy for Rei which is optimal, and given this, it is not too hard to determine the expected score when both players play optimally. Unfortunately, it turns out that Rei does not have a "simple" optimal strategy in the semi-restricted D-game for almost every Eulerian tournament D, see Theorem 18 for a precise statement. Despite this significant hurdle, we can still determine the expected score when both players play optimally.

Theorem 4. *If D is a non-empty Eulerian tournament, then*

$$S_D(n \cdot \mathbf{1}) = \Theta_D(\sqrt{n}).$$

In fact, we will show more generally that $S_D(n \cdot \mathbf{1}) = \Theta_D(\sqrt{n})$ whenever D is an Eulerian digraph satisfying certain spectral conditions, see Theorem 11.

Lastly, it is natural to consider not just semi-restricted games where one player is restricted, but also games where both players are restricted. Indeed, several games of this form appeared in the manga Tobaku Mokushiroku Kaiji [4], and one such game was studied by Spiro [11]. It turns out that these games are much simpler to analyze. In particular, in every such game, it is optimal for both players to uniformly at random choose from their multiset of available actions each round, see Appendix C for more on this.

1.3 Organization and Notation

The remainder of this paper is organized as follows. In Section 2 we provide a more formal definition of $S_D(r)$ and determine the optimal strategies for every semi-restricted D-game when D has at most 3 vertices, in particular proving the first half of Theorem 1. In Section 3 we prove general bounds on $S_D(r)$ and prove Theorems 2, 3, 4 and the second half of Theorem 1. In Section 4 we show that almost every (Eulerian) tournament does not have a "simple" optimal strategy for Rei. We end with some open problems in Section 5.

We recall the following standard asymptotic notation. For two functions f, g depending on n, we write $g = O(f)$ to mean there exists some $C > 0$ such that $g(n) \leqslant Cf(n)$ for all n. Similarly $g = \Omega(f)$ indicates $g(n) \geqslant cf(n)$ for some $c > 0$, and $g = \Theta(f)$ means $g = O(f)$ and $g = \Omega(f)$. We write, for example, $g = O_D(f)$ to indicate that there is

some $C > 0$ depending on the parameter D such that $g(n) \leqslant Cf(n)$ for all n. We write $g = o(f)$ to indicate $\lim_{n\to\infty} \frac{g(n)}{f(n)} = 0$.

Given a vector r, we define the *support* $\mathrm{supp}(r)$ to be the set of v with $r_v > 0$. Similarly given a random variable X, we define its *support* $\mathrm{supp}(X)$ to be the set of x with $\Pr(X = x) > 0$. Given an event A we let $\mathbb{1}(A)$ denote the indicator function which is 1 if A occurs and 0 otherwise.

2 Optimal Strategies

In this section we prove several results about optimal strategies. We begin by establishing formal definitions for some of the terms that were informally defined in the introduction.

Given a digraph D, a *strategy for Norman* in the semi-restricted D-game is a function $\mathcal{N}$ from restriction vectors to random variables such that $\mathrm{supp}(\mathcal{N}(r)) \subseteq V(D)$ for all r. A *strategy for Rei* is a function $\mathcal{R}$ from restriction vectors to random variables such that $\mathrm{supp}(\mathcal{R}(r)) \subseteq \mathrm{supp}(r)$ for all r. Informally, this simply says that Norman is allowed to play any vertex of D, while Rei can only play vertices in $\mathrm{supp}(r)$.

We let δ_v be the vector which has a 1 in the position corresponding to $v \in V(D)$ and 0's everywhere else. Given strategies $\mathcal{N}, \mathcal{R}$ for Norman and Rei respectively and a restriction vector r, we define the *score* (for Norman) $\mathbf{S}_D(r; \mathcal{N}, \mathcal{R})$ recursively by setting $\mathbf{S}_D(r; \mathcal{N}, \mathcal{R}) = 0$ if $r = (0, \ldots, 0)$, and otherwise having

$$\mathbf{S}_D(r; \mathcal{N}, \mathcal{R}) = \begin{cases} 1 + \mathbf{S}_D(r - \delta_{\mathcal{R}(r)}; \mathcal{N}, \mathcal{R}) & \text{if } \mathcal{R}(r) \in N^+(\mathcal{N}(r)), \\ -1 + \mathbf{S}_D(r - \delta_{\mathcal{R}(r)}; \mathcal{N}, \mathcal{R}) & \text{if } \mathcal{R}(r) \in N^-(\mathcal{N}(r)), \\ \mathbf{S}_D(r - \delta_{\mathcal{R}(r)}; \mathcal{N}, \mathcal{R}) & \text{otherwise.} \end{cases}$$

Informally, this says Norman and Rei play a round of the semi-restricted D-game with parameter r, record the change in score based on whether Norman's choice beats Rei's or not, and then the game continues with one action of $\mathcal{R}(r)$ removed for Rei. Note that $\mathbf{S}_D(r; \mathcal{N}, \mathcal{R})$ is a random variable. We define *the expected score* $S_D(r; \mathcal{N}, \mathcal{R}) := \mathbb{E}[\mathbf{S}_D(r; \mathcal{N}, \mathcal{R})]$.

For r a restriction vector and $\mathcal{R}$ a Rei's strategy, we define

$$S_D(r; \mathcal{R}) = \max_{\mathcal{N}} S_D(r; \mathcal{N}, \mathcal{R}) \qquad \text{and} \qquad S_D(r) = \min_{\mathcal{R}} S_D(r; \mathcal{R}),$$

where the maximum and minimums run through all possible strategies of Norman and Rei as appropriate. 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. Intuitively, $S_D(r; \mathcal{R})$ is Norman's score if he plays optimally and Rei uses strategy $\mathcal{R}$ and has restriction vector r. If $\mathcal{R}$ is such that $S_D(r) = S_D(r; \mathcal{R})$ for all vectors r, then we say that $\mathcal{R}$ is an *optimal strategy* for Rei[2].

We start with a basic fact about the behavior of S_D.

[2] As defined, it is not immediately clear that Norman can achieve a score of $S_D(r)$ without knowing that Rei is using some given optimal strategy. But this follows from von Neuman's famed minimax theorem for finite two player zero-sum games.

Fact 5. *Let $\mathcal{R}$ be a strategy for Rei. If r is a restriction vector and $p_u = \Pr(\mathcal{R}(r) = u)$ for each $u \in V(D)$, then*

$$S_D(r;\mathcal{R}) = \max_v \left\{ \sum_{u \in N^+(v)} p_u - \sum_{u \in N^-(v)} p_u \right\} + \sum_u p_u S_D(r - \delta_u;\mathcal{R}).$$

Intuitively, this result says that if Rei is playing according to strategy $\mathcal{R}$, then Norman should always play the vertex v which maximizes his expected score for each round, i.e. the v such that $\sum_{u \in N^+(v)} p_u - \sum_{u \in N^-(v)} p_u$ is as large as possible. A close analog of this is proven in Spiro [11, Lemma 2.1]. One can easily give a formal proof of this results by adjusting the corresponding proof in [11].

The key lemma we need for this section is the following. Roughly speaking, it says that in any semi-restricted D-game, the optimal expected score will not change dramatically if among all the remaining options of Rei, one action of u is changed into one action of v.

Lemma 6. *Given $u, v \in V(D)$ and a restriction vector r with $r_u, r_v \geqslant 1$, we have*

$$S_D(r - \delta_u) \leqslant S_D(r - \delta_v) + \alpha(u, v), \tag{1}$$

where

$$\alpha(u,v) := \begin{cases} 2 & \text{if } N^+(u) \cap N^-(v) \neq \emptyset, \\ 0 & \text{if } N^+(u) \cup N^-(v) = \emptyset, \\ 1 & \text{otherwise.} \end{cases}$$

Moreover, we have strict inequality in (1) *when either $N^-(u) \neq \emptyset$ or $\alpha(u, v) = 2$.*

We note that this inequality can be sharp, for example if u is an isolated vertex and $N^-(v) = \emptyset$, but typically it will not be (e.g. if $N^-(u) \neq \emptyset$ or $\alpha(u, v) = 2$).

Proof. We first assume the condition $N^-(u) \neq \emptyset$ and prove the strict inequality

$$S_D(r - \delta_u) < S_D(r - \delta_v) + \alpha(u, v).$$

After its proof we briefly discuss the other scenarios of this theorem, i.e. when $\alpha(u, v) = 2$ or when we conclude non-strict inequalities.

Suppose for contradiction that there is some r with $S_D(r - \delta_u) \geqslant S_D(r - \delta_v) + \alpha(u, v)$, and choose this r with $\sum r_w$ as small as possible. Let $\mathcal{R}$ be an optimal strategy for Rei, and let $p_w = \Pr(\mathcal{R}(r - \delta_v) = w)$. The idea now is for Rei to consider a different (and possibly suboptimal strategy) $\mathcal{R}'$ which has Rei acting as if her restriction vector is $r - \delta_v$ whenever it is $r - \delta_u$. That is, we would like to choose $\mathcal{R}'$ such that $\Pr(\mathcal{R}'(r - \delta_u) = w) = p_w$ for all w. However, this can cause issues when $r_u = 1$, since in this case $\mathcal{R}'$ may choose u with positive probability despite $r - \delta_u$ having no copies of u available for her to play. We get around this by having $\mathcal{R}'$ select v whenever $\mathcal{R}$ would have selected u. More precisely, we define a strategy $\mathcal{R}'$ by having $\mathcal{R}'(r') = \mathcal{R}(r')$ if $r' \neq r - \delta_u$, and otherwise setting

$$p'_w := \Pr(\mathcal{R}'(r-\delta_u)=w) = \begin{cases} p_u+p_v & w=v, \\ 0 & w=u, \\ p_w & w\neq u,v. \end{cases}$$

Observe that $\mathcal{R}'$ is always a strategy for Rei (i.e. $\mathcal{R}'(r')\subseteq \operatorname{supp}(r')$ for all r') since $r_v>0$ and $\mathcal{R}$ is a strategy for Rei.

First assume $p_u=1$, which means the optimal strategy $\mathcal{R}$ always has Rei choosing u at $r-\delta_v$. By Fact 5, our assumption $N^-(u)\neq\emptyset$, and the fact D has no self-loops, we have

$$S_D(r-\delta_v) = 1+S_D(r-\delta_u-\delta_v). \tag{2}$$

The possibly suboptimal strategy $\mathcal{R}'$ always has Rei choosing v at $r-\delta_u$, so Fact 5 implies

$$S_D(r-\delta_u;\mathcal{R}') = \begin{cases} 1+S_D(r-\delta_u-\delta_v), & \text{if } N^-(v)\neq\emptyset \\ S_D(r-\delta_u-\delta_v), & \text{if } N^-(v)=\emptyset \end{cases}. \tag{3}$$

Hence, if $N^-(v)=\emptyset$, we have

$$S_D(r-\delta_u)\leqslant S_D(r-\delta_u;\mathcal{R}')<S_D(r-\delta_v)\leqslant S_D(r-\delta_v)+\alpha(u,v),$$

where the first inequality used that $\mathcal{R}'$ is possibly suboptimal, the second used (2) and (3), and the last used that $\alpha(u,v)\geqslant 0$ always. If $N^-(v)\neq\emptyset$, we have

$$S_D(r-\delta_u)\leqslant S_D(r-\delta_u;\mathcal{R}')=S_D(r-\delta_v)<S_D(r-\delta_v)+\alpha(u,v),$$

since in this case $\alpha(u,v)>0$. Either way contradicts our assumption, so we must have $p_u\neq 1$.

Define $e_w=\sum_{w'\in N^+(w)}p_{w'}-\sum_{w'\in N^-(w)}p_{w'}$, which is the expected score when one round is played at $r-\delta_v$ when Rei follows $\mathcal{R}$ and Norman picks w. By Fact 5 and $\mathcal{R}$ being optimal, we have

$$S_D(r-\delta_v)=\max_w e_w+\sum_w p_w S_D(r-\delta_v-\delta_w). \tag{4}$$

Similarly, we define $e'_w=\sum_{w'\in N^+(w)}p'_{w'}-\sum_{w'\in N^-(w)}p'_{w'}$. By definition of p'_w (which differs by either 0 or $\pm p_u$ from p_w), we have

$$\begin{aligned} e'_w-e_w &= \Big[\sum_{w'\in N^+(w)}p'_{w'}-p_{w'}\Big]-\Big[\sum_{w'\in N^-(w)}p'_{w'}-p_{w'}\Big] \\ &= \big[p_u\mathbb{1}(v\in N^+(w))-p_u\mathbb{1}(u\in N^+(w))\big]-\big[p_u\mathbb{1}(v\in N^-(w))-p_u\mathbb{1}(u\in N^-(w))\big] \\ &\leqslant p_u\mathbb{1}(v\in N^+(w))+p_u\mathbb{1}(u\in N^-(w)) \\ &= p_u\mathbb{1}(w\in N^-(v))+p_u\mathbb{1}(w\in N^+(u))\leqslant p_u\alpha(u,v), \end{aligned}$$

where this last step uses the definition of $\alpha(u,v)$. Using this together with Fact 5 gives

$$S_D(r-\delta_u;\mathcal{R}') \leqslant \max_w e_w + p_u\alpha(u,v) + \sum_{w\neq u,v} p_w S_D(r-\delta_u-\delta_w) + (p_u+p_v)S_D(r-\delta_u-\delta_v), \quad (5)$$

where implicitly we used that $S_D(r-\delta_u-\delta_w;\mathcal{R}') = S_D(r-\delta_u-\delta_w)$ since $\mathcal{R}'$ agrees with the optimal strategy $\mathcal{R}$ for every vector except $r-\delta_u$. Since $S_D(r-\delta_u) \leqslant S_D(r-\delta_u;\mathcal{R}')$, we can combine (4) and (5) and obtain

$$\begin{aligned} S_D(r-\delta_u) - S_D(r-\delta_v) &\leqslant S_D(r-\delta_u;\mathcal{R}') - S_D(r-\delta_v) \\ &\leqslant p_u\alpha(u,v) + \sum_{w\neq u} p_w(S_D(r-\delta_u-\delta_w) - S_D(r-\delta_v-\delta_w)) \\ &< p_u\alpha(u,v) + \sum_{w\neq u} p_w\alpha(u,v) \qquad (6) \\ &= \alpha(u,v), \end{aligned}$$

where the inequality (6) used the assumptions that $p_u \neq 1$ (so $p_w > 0$ for some $w \neq u$) and that r was a minimum counterexample to the desired strict inequality. As this contradicts our assumption, we conclude our desired strict inequality.

This completes the proof in the case $N^-(u) \neq \emptyset$. If $N^-(u) = \emptyset$ but $\alpha(u,v) = 2$, the condition $p_u = 1$ implies the following identity, which is slightly different from (2),

$$S_D(r-\delta_v) = S_D(r-\delta_v-\delta_u).$$

The identity (3) still holds for $p_u = 1$. Together they imply $|S_D(r-\delta_u) - S_D(r-\delta_v)| \leqslant 1 < \alpha(u,v)$. So $p_u = 1$ still contradicts the assumption for r, and when $p_u \neq 1$ we can argue as above and conclude (1) with a strict inequality. If we are in the situation where $N^-(u) = \emptyset$ and $\alpha(u,v) \neq 2$, then the same argument as above gives (6) with a non-strict inequality, giving the desired result. □

We can use Lemma 6 to establish the optimal strategies for Rei for every digraph on at most 3 vertices. Most of these cases are trivial. Indeed, recall that a vertex v in a digraph is called a *source* if $N^-(v) = \emptyset$. If there is a source v such that $N^+(v)$ contains all of the non-source vertices, then an optimal strategy for Norman is to always choose v (which will allow him to win every round that Rei does not play a source, and Norman can only draw when Rei plays a source). Facing this strategy of Norman, every strategy of Rei gives the same outcome, hence is trivially optimal. One can check that such a source exists in every digraph with at most 3 vertices except for the circuit on 3 vertices and the directed path on 3 vertices, so it remains to address these two cases.

Recall that if D is the directed cycle $1 \to 2 \to 3 \to 1$, then the D-game is just RPS. In Section 1 we mentioned the following greedy strategy $\mathcal{R}_g$ for Rei in this semi-restricted D-game: When she still has all three options available, she chooses each with uniform probability; When she has two options, she chooses the stronger one with probability 2/3 and the weaker one with probability 1/3; When she has only one option, she chooses this with probability 1. The next theorem implies the first half of Theorem 1.

Theorem 7. *The greedy strategy $\mathcal{R}_g$ is the unique optimal strategy for Rei in the semi-restricted D-game when D is the directed cycle $1 \to 2 \to 3 \to 1$.*

Proof. Let $\mathcal{R}$ be an optimal strategy for Rei. We shall prove that $\Pr(\mathcal{R}(r) = u) = \Pr(\mathcal{R}_g(r) = u)$ for all vertices u and vectors r by induction on the quantity $r_1+r_2+r_3$. The base case when $r_1 + r_2 + r_3 = 0$ is trivial. Now suppose $\Pr(\mathcal{R}(r') = u) = \Pr(\mathcal{R}_g(r') = u)$ for all vertices u and vectosr r' satisfying $r'_1+r'_2+r'_3 < r_1+r_2+r_3$. Let $p_u := \Pr(\mathcal{R}(r) = u)$ for each $u \in \{1, 2, 3\}$.

When $|\mathrm{supp}(r)| = 1$, by the definition of a strategy for Rei, we have $\Pr(\mathcal{R}(r) = u) = \Pr(\mathcal{R}_g(r) = u)$ for all u.

When $|\mathrm{supp}(r)| = 2$, we can assume without loss of generality that $r_1, r_2 > 0$. Suppose $p_1 = 2/3 + \epsilon$ with $\epsilon > 0$, which means $p_2 = 1/3 - \epsilon$. Applying Fact 5 to both $\mathcal{R}$ and $\mathcal{R}_g$ gives

$$S_D(r; \mathcal{R}) = \left(\frac{1}{3} + 2\epsilon\right) + \left(\frac{2}{3} + \epsilon\right) S_D(r - \delta_1) + \left(\frac{1}{3} - \epsilon\right) S_D(r - \delta_2),$$
$$S_D(r; \mathcal{R}_g) = \frac{1}{3} + \frac{2}{3} S_D(r - \delta_1) + \frac{1}{3} S_D(r - \delta_2),$$

where here we used the assumptions that $\mathcal{R}$ is optimal and that $\mathcal{R}_g$ coincides with $\mathcal{R}$ on $r' = r - \delta_u$ for $u = 1, 2$. These two identities imply

$$S_D(r; \mathcal{R}_g) - S_D(r; \mathcal{R}) = \epsilon(S_D(r - \delta_2) - S_D(r - \delta_1) - 2) < 0,$$

where the inequality is by Lemma 6 with $\alpha(2, 1) = 2$ (since $N^+(2) \cap N^-(1) = \{3\}$). This contradicts $\mathcal{R}$ being an optimal strategy, so we can not have $p_1 = 2/3 + \epsilon$ for any $\epsilon > 0$. If $p_1 = 2/3 - \epsilon$ for some $\epsilon > 0$, we now have

$$S_D(r; \mathcal{R}) = \left(\frac{1}{3} + \epsilon\right) + \left(\frac{2}{3} + \epsilon\right) S_D(r - \delta_1) + \left(\frac{1}{3} - \epsilon\right) S_D(r - \delta_2).$$

Similar to before we find

$$S_D(r; \mathcal{R}_g) - S_D(r; \mathcal{R}) = \epsilon(S_D(r - \delta_1) - S_D(r - \delta_2) - 1) < 0,$$

where the inequality is again by Lemma 6. This is also a contradiction, so we must have $p_1 = 2/3$ and $p_2 = 1/3$, which means $\Pr(\mathcal{R}(r) = u) = \Pr(\mathcal{R}_g(r) = u)$ for all u as desired.

When $|\mathrm{supp}(r)| = 3$, we let $\Delta := \max\{p_2 - p_3, p_3 - p_1, p_1 - p_2\}$. Note that $\Delta \geqslant 0$. By Fact 5 and the assumption of $\mathcal{R}$ being optimal, we have

$$S_D(r; \mathcal{R}) = \Delta + p_1 S_D(r - \delta_1) + p_2 S_D(r - \delta_2) + p_3 S_D(r - \delta_3).$$

Our proof now divides into three cases.

Case 1. Exactly one term in $\{p_2 - p_3, p_3 - p_1, p_1 - p_2\}$ equals Δ.

Without loss of generality, we can assume $p_2 - p_3 = \Delta$. Take

$$\epsilon = \frac{1}{3} \min\{|\Delta - (p_3 - p_1)|, |\Delta - (p_1 - p_2)|\} > 0.$$

Consider a different strategy $\mathcal{R}'$ for Rei with

$$\Pr(\mathcal{R}'(r) = 1) = p_1,\ \Pr(\mathcal{R}'(r) = 2) = p_2 - \epsilon,\ \Pr(\mathcal{R}'(r) = 3) = p_3 + \epsilon,$$

and $\Pr(\mathcal{R}'(r') = u) = \Pr(\mathcal{R}(r') = u)$ for all vertices u and vectors $r' \neq r$. By Fact 5 and our choice of ϵ, we have

$$S_D(r; \mathcal{R}') = (\Delta - 2\epsilon) + p_1 S_D(r - \delta_1) + (p_2 - \epsilon) S_D(r - \delta_2) + (p_3 + \epsilon) S_D(r - \delta_3).$$

Thus

$$S_D(r; \mathcal{R}') - S_D(r; \mathcal{R}) = \epsilon(S_D(r - \delta_3) - S_D(r - \delta_2) - 2) < 0,$$

where the inequality is by Lemma 6. This contradicts $\mathcal{R}$ being optimal, so Case 1 can not happen.

Case 2. Exactly two expressions in $\{p_2 - p_3, p_3 - p_1, p_1 - p_2\}$ equal Δ.

Without loss of generality, we assume $p_2 - p_3 = p_3 - p_1 = \Delta$. This implies $p_3 = 1/3$, $p_2 = 1/3 + \Delta$, and $p_1 = 1/3 - \Delta$. Note that $\Delta > 0$, as otherwise $p_1 - p_2$ also attains Δ. Similar to the above arguments, we can compute

$$S_D(r; \mathcal{R}_g) - S_D(r; \mathcal{R}) = \Delta(S_D(r - \delta_1) - S_D(r - \delta_2) - 1) < 0,$$

where the inequality is by Lemma 6. This gives a contradiction, so Case 2 can not happen.

Case 3. Every expression in $\{p_2 - p_3, p_3 - p_1, p_1 - p_2\}$ equals Δ.

We have $p_2 - p_3 = p_3 - p_1 = p_1 - p_2$, which implies $p_1 = p_2 = p_3 = 1/3$ and hence $\Pr(\mathcal{R}(r) = u) = \Pr(\mathcal{R}_g(r) = u)$ for all u. As neither of the previous two cases happen, we must be in this case, so we obtain the desired conclusion. □

The only remaining digraph on 3 vertices is the directed path, and in this case it turns out there are many optimal strategies for Rei.

Theorem 8. *The optimal strategies for Rei in the semi-restricted D-game when D is the directed path $1 \to 2 \to 3$ are exactly those $\mathcal{R}$ satisfying $\Pr(\mathcal{R}(r) = 3) = 1/2$ whenever $\{3\} \subsetneq \mathrm{supp}(r)$.*

That is, Rei will always play 3 with probability 1/2 provided she can play 3 and at least one other option. The intuition for Theorem 8 is as follows: in any particular round, suppose Rei plays v with probability p_v. It is clear that Norman will only ever play 1 or 2 for his turn, and for a given round these gives him expected payoffs of p_2 and $p_1 - p_3$, respectively. These two payoffs are equal precisely when $p_3 = 1/2$. Thus the strategies in Theorem 8 are exactly those such that Norman is indifferent between which (reasonable) vertex he should play. The proof of Theorem 8 is similar to that of Theorem 7, and as such we have relegated its proof to Appendix A.

3 Bounding the Score

3.1 General bounds

Our proofs will need the following technical result, which will be used to control the expected number of rounds remaining in the semi-restricted D-game after one option has been depleted.

Lemma 9. *Let $p \in (0,1]$ and N be an integer. Let $X = X_1 + \cdots + X_N$ where the X_i are iid random geometric variables with parameter p. Then*

$$\mathbb{E}[(\mathbb{E}X - X) \cdot \mathbb{1}(X \leqslant \mathbb{E}X)] = O(\sqrt{N}/p)$$

Proof. Note that $\mathrm{Var}(X) = N(1-p)/p^2$. Therefore by Cauchy-Schwarz inequality

$$\mathbb{E}[(\mathbb{E}X - X)\mathbb{1}(X \leqslant \mathbb{E}X)] \leqslant O(\sqrt{\mathrm{Var}(X)\mathrm{Var}(\mathbb{1}(X \leqslant \mathbb{E}X))} = O(\sqrt{N}/p),$$

where we use the fact that $\mathrm{Var}(\mathbb{1}(X \leqslant \mathbb{E}X)) = O(1)$. □

We can use this technical result to prove Theorem 2.

Proof of Theorem 2. Recall that we wish to show

$$S_D(n \cdot \mathbf{1}) = \max_v \{d^+(v) - d^-(v)\}n + O_D(\sqrt{n}).$$

The lower bound $S_D(n\cdot\mathbf{1}) \geqslant \max_v\{d^+(v)-d^-(v)\}n$ follows form Norman deterministically choosing v every round. It remains to exhibit a strategy of Rei such that Norman can gain at most $\max_v\{d^+(v)-d^-(v)\}n+O_D(\sqrt{n})$ points in expectation against this strategy.

Intuitively, Rei's strategy will be to play uniformly at random until some option is depleted, after which she plays arbitrarily. To more formally analyze this strategy, we will have Rei generate an infinite random string $\pi = \pi_1\pi_2\cdots$ where each π_t is chosen uniformly and independently amongst $V(D)$. In the t-th round, Rei will play π_t, unless that option has been depleted, in which case she plays arbitrarily.

Let $k = |V(D)|$. For $w \in V(D)$, let T_w denote the smallest integer t such that $|\{s \leqslant t : \pi_s = w\}| = n$, and let $T = \min_w T_w$. Note that before the T-th round of the game, Rei can still play every option (and hence does so uniformly), and that some option is depleted after completion of the T-th round. Thus for $t \leqslant T$, if Norman plays v in the t-th round, then the expected increase in score for Norman is exactly $k^{-1}(d^+(v) - d^-(v))$, so his maximum expected increase in score is at most $k^{-1}\max_v\{d^+(v) - d^-(v)\}$. Using this and that Norman's score increases by at most 1 for each round after the T-th, we find

$$\begin{aligned}\mathbb{E}[S_D(r)|T] &\leqslant T \cdot k^{-1} \max_v\{d^+(v) - d^-(v)\} + (kn - T) \\ &\leqslant \max_v\{d^+(v) - d^-(v)\}n + (kn - T), \end{aligned} \tag{7}$$

where this last step used that deterministically $T \leqslant kn$. It remains to upper bound the expected value of $kn - T$.

Observe that T_w has the distribution of the sum of n independent geometric random variables with parameter $1/k$. Thus $\mathbb{E}T_w = kn$, and Lemma 9 gives

$$\mathbb{E}[(kn - T_w)\mathbb{1}(T_w \leqslant kn)] = O_D(\sqrt{n}).$$

It follows that

$$\mathbb{E}[kn - T] = \mathbb{E}[\max_w\{kn - T_w\}] \leqslant \sum_w \mathbb{E}[(kn - T_w)\mathbb{1}(T_w \leqslant kn)] = O_D(\sqrt{n}).$$

Combining this with (7) and using the tower property of conditional expectation gives the desired result. □

A similar approach gives Theorem 3.

Proof of Theorem 3. Recall that we wish to prove

$$S_D(r) \leqslant \max_v \left\{ \sum_{u \in N^+(v)} r_u - \sum_{u \in N^-(v)} r_u \right\} + O_D(M^{2/3}),$$

where $M = \max_v r_v$. The approach we use is very similar to that of Theorem 2, so we will omit some of the redundant details. Intuitively, we will prove the lower bound by having Rei play each option w with probability $r_w / \sum_u r_u$ until an option is depleted, after which she plays arbitrarily. However, we will first need to "trim" r to ignore vertices with r_v small.

To be more precise, let r' be the vector defined by $r'_w = r_w$ if $r_w \geqslant M^{2/3}$ and $r'_w = 0$ otherwise. Rei generates an infinite random string $\pi = \pi_1\pi_2\cdots$ where $\pi_t = w$ with probability $r'_w / \sum_u r'_u$ for all $w \in \mathrm{supp}(r')$ independently of every other π_s. In the t-th round Rei will play π_t unless that option is depleted, in which case she plays arbitrarily.

For $w \in \mathrm{supp}(r')$, let T_w denote the smallest integer t such that $|\{s \leqslant t : \pi_s = w\}| = r'_w$, and let $T = \min_{w \in \mathrm{supp}(r')} T_w$. Note that for $t \leqslant T$, if Norman plays v in the t-th round then the expected increase in score for Norman is exactly $(\sum_u r'_u)^{-1}(\sum_{u \in N^+(v)} r'_u - \sum_{u \in N^-(v)} r'_u)$. Using this, that Norman's score increases by at most 1 each round, and that $T \leqslant \sum_u r'_u$, we see that

$$\mathbb{E}[S_D(r)|T] \leqslant \max_v \left\{ \sum_{u \in N^+(v)} r'_u - \sum_{u \in N^-(v)} r'_u \right\} + \left(\sum_u r_u - T \right). \tag{8}$$

Observe that each T_w random variable has the distribution of the sum of $r'_w = r_w$ independent geometric random variables with parameter $r_w / \sum_u r'_u$. Thus $\mathbb{E}T_w = \sum_u r'_u$, and Lemma 9 gives

$$\mathbb{E}[(\sum_u r'_u - T_w)\mathbb{1}(T_w \leqslant \sum_u r'_u)] = O_D(r_w^{-1/2} \sum_u r'_u) = O_D(M^{2/3}),$$

where this last step used $r_w \geqslant M^{2/3}$ and $\sum_u r'_u \leqslant k \cdot M$. It follows that

$$\mathbb{E}[\sum_u r'_u - T] = \mathbb{E}[\max_w \{\sum_u r'_u - T_w\}] \leqslant \sum_w \mathbb{E}[(\sum_u r'_u - T_w)\mathbb{1}(T_w \leqslant \sum_u r'_u)] = O_D(M^{2/3}).$$

Combining this with (8), the tower property of conditional expectation, and that $\sum_u r_u - \sum_u r'_u = O_D(M^{2/3})$ gives the desired result. □

3.2 Spectral bounds

As we saw in the case of RPS, there are semi-restricted games where Rei's optimal strategy is to start by playing uniformly at random until some option runs out. In this section, we will classify a class of graphs for which Rei's optimal strategy is to play approximately uniformly until some option runs out. We do so by identifying a property which implies that the expected loss at any given step is proportional to its deviation from the uniform strategy.

Given a digraph D, we identify $V(D)$ with $[k]$ for some integer k and let A_D be the *skew adjacency matrix* of D, i.e.

$$(A_D)_{ij} = \begin{cases} 1, & \text{if } ij \in E(D) \\ -1 & \text{if } ji \in E(D) \\ 0 & \text{otherwise} \end{cases}.$$

Our main focus for this subsection will be digraphs satisfying

$$\mathrm{Null}(A_D) = \mathrm{span}(\mathbf{1}).$$

To motivate this, observe that at any given state r, if p is the probability vector such that Rei plays vertex u with probability p_u, then Rei's expected loss for this round will be 0 if and only if $A_D p = 0$. Therefore, for digraphs with $\mathrm{Null}(A_D) = \mathrm{span}(\mathbf{1})$, the steps where Rei's expected loss is 0 are the ones where Rei picks her option uniformly at random. We will now show that in fact the expected loss for such digraphs is always at least proportional to the maximum deviation between p_v and $1/k$.

Lemma 10. *For any k-vertex digraph D with $\mathrm{Null}(A_D) = \mathrm{span}(\mathbf{1})$, there exists a constant $\alpha_D > 0$ such that for any probability vector p indexed by $V(D)$, we have*

$$\max_v \left\{ \sum_{u \in N^+(v)} p_u - \sum_{u \in N^-(v)} p_u \right\} \geqslant \alpha_D \max_v |p_v - 1/k|.$$

Proof. First notice that the left-hand side of our inequality satisfies

$$\max_v \left\{ \sum_{u \in N^+(v)} p_u - \sum_{u \in N^-(v)} p_u \right\} = \max_v (A_D p)_v \geqslant k^{-1} \max_v |(A_D p)_v|, \tag{9}$$

where the last inequality holds because $\sum_v (A_D p)_v = 0$ and because $x = (x_1, \ldots, x_k)$ with $\sum_i^k x_i = 0$ and $x_1 \geqslant \cdots \geqslant x_k$ satisfies

$$kx_1 \geqslant \sum_{i:\ x_i \geqslant 0} x_i = \sum_{i:\ x_i \leqslant 0} |x_i| \geqslant |x_k|.$$

Since A_D is skew-symmetric and $\dim(\mathrm{Null}(A_D)) = 1$, it admits an eigendecomposition $A_D = \sum_{i=1}^k \lambda_i e_i e_i^*$ with $0 = |\lambda_1| < |\lambda_2| \leqslant \cdots \leqslant |\lambda_k|$, $e_i^* e_i = 1$, and $e_i^* e_j = 0$ for all $i \neq j$. Notice that $A_D p = A_D(p - k^{-1} \cdot \mathbf{1})$, we have

$$\begin{aligned} ||A_D p||_2 = \left|\left|A_D\left(p - k^{-1}\cdot \mathbf{1}\right)\right|\right|_2 &= \left|\left|\sum_{i=2}^k e_i \left[\lambda_i e_i^*\left(p - k^{-1}\cdot \mathbf{1}\right)\right]\right|\right|_2 \\ &= \sqrt{\sum_{i=2}^k |\lambda_i|^2 \left|e_i^*\left(p - k^{-1}\cdot \mathbf{1}\right)\right|^2} \geqslant |\lambda_2| \cdot \left|\left|p - k^{-1}\cdot \mathbf{1}\right|\right|_2 \geqslant |\lambda_2| \max_v |p_v - 1/k|\,. \end{aligned}$$

Hence

$$\max_v |(A_D p)_v| \geqslant \frac{1}{k}\,||A_D p||_2 \geqslant \frac{|\lambda_2|}{k} \max_v |p_v - 1/k|\,,$$

so by (9) the statement follows with $\alpha_D = |\lambda_2|/k^2$. □

Using this we can prove the following lower bound on $S_D(n \cdot \mathbf{1})$ whenever $\mathrm{Null}(A_D) = \mathrm{span}(\mathbf{1})$.

Theorem 11. *If D is a digraph on at least three vertices with $\mathrm{Null}(A_D) = \mathrm{span}(\mathbf{1})$, then*

$$S_D(n \cdot \mathbf{1}) = \Omega(\sqrt{n}).$$

Proof. We identify $V(D)$ with $[k]$ for $k \geqslant 3$ and let $N = kn$ be the total number of rounds in this game. Fix some optimal strategy $\mathcal{R}$ for Rei. We divide the game into two phases: the first phase is when every option is still available to Rei, and the second phase is when at least one option has been depleted. Let T be the (random) number of steps in the first phase.

Let $\mathbf{X} = (X_1, X_2, \ldots, X_N)$ be the stochastic process where X_i denotes Rei's choice when $i \leqslant T$, and for $i > T$ the X_i are independently uniformly random over $[k]$. We define

$$p_v^t = \Pr(X_t = v | X_1, X_2, \ldots, X_{t-1}).$$

For $t \leqslant T$, p_v^t corresponds to the probability of Rei choosing v at the t-th round given her previous choices.

We consider the following greedy strategy for Norman: Suppose at the t-th round, the restriction vector for Rei is r with $p_v = \Pr(\mathcal{R}(r) = v)$ for all v, then Norman will deterministically pick w that maximizes his expected gain, i.e.

$$w = \arg\max_v \left\{ \sum_{u \in N^+(v)} p_u - \sum_{u \in N^-(v)} p_u \right\}.$$

Here we break ties arbitrarily. If $t \leqslant T$, we have $p_v = p_v^t$. Then, by Lemma 10, Norman's expected gain at this round is

$$\max_v \left\{ \sum_{u \in N^+(v)} p_u - \sum_{u \in N^-(v)} p_u \right\} = \max_v \left\{ \sum_{u \in N^+(v)} p_u^t - \sum_{u \in N^-(v)} p_u^t \right\} \geqslant \alpha_D |p_1^t - 1/k|.$$

If $t > T$, there exists some v with $p_v = 0$. Then, by Lemma 10, Norman's expected gain at this round is

$$\max_v \left\{ \sum_{u \in N^+(v)} p_u - \sum_{u \in N^-(v)} p_u \right\} \geqslant \alpha_D / k$$

We let S_1 denote Norman's expected score during the first phase under this strategy and S_2 his score during the second phase, so that his total expected score is $S_1 + S_2$. By above analysis, we have

$$S_1 \geqslant \mathbb{E}\left[\sum_{t=1}^{T} \alpha_D |p_1^t - 1/k| \right] = \alpha_D \mathbb{E}\left[\sum_{t=1}^{N} |p_1^t - 1/k| \right], \tag{10}$$

$$S_2 \geqslant \mathbb{E}\left[\sum_{t=T+1}^{N} \alpha_D / k \right] = \frac{\alpha_D}{k} \cdot \mathbb{E}[N - T]. \tag{11}$$

Here, the second identity of (10) follows from $p_1^t = 1/k$ when $t > T$.

From (10) and (11), we see that $S_1, S_2 \geqslant 0$, and thus to prove the result it suffices to show $\max\{\mathbb{E}\sum_t |p_1^t - 1/k|, \mathbb{E}[N - T]\} = \Omega(\sqrt{n})$. To this end, fix some small constant $c > 0$ to be chosen later. We may assume from now on that

$$\mathbb{E}\left[\sum_{t=1}^{N} |p_1^t - 1/k| \right] \leqslant c\sqrt{n}, \tag{12}$$

as otherwise we are done by (10).

We claim that, if c is sufficiently small, then with probability at least $1/8$ the vertex 1 appears at least $n + \Omega(\sqrt{n})$ times in $\mathbf{X}$. This will imply our theorem. Indeed, if 1 appears $n + Y$ times in $\mathbf{X}$, then necessarily $N - T \geqslant Y$ since $(X_1, \ldots, X_{T-1})$ by definition has each symbol appearing less than n times. Thus this claim together with (11) implies $S_2 = \Omega(\sqrt{n})$, giving the result. It remains to prove the claim.

Let $f(\mathbf{X})$ be the number of times 1 occur in $\mathbf{X}$. We will show that we can couple $f(\mathbf{X})$ with a binomial random variable $\mathrm{Bin}(N, 1/k)$ such that

$$\mathbb{E}[|f(\mathbf{X}) - \mathrm{Bin}(N, 1/k)|] \leqslant \mathbb{E}\left[\sum_{t=1}^{N} |p_1^t - 1/k| \right].$$

To do so, we consider independent uniform random variables $U_1, U_2, \ldots, U_N \sim U[0, 1]$. Let $Y_t = \mathbb{1}(U_t \leqslant 1/k)$. Notice that

$$\mathrm{Bin}(N, 1/k) \sim \sum_{1 \leqslant t \leqslant N} Y_t.$$

To couple X_t with Y_t, we let

$$X_t = j \quad \text{if} \quad U_t \in \left(\sum_{i<j} p_i^t, \sum_{i \leqslant j} p_i^t \right].$$

With this coupling,

$$\begin{aligned}
\Pr(\mathbb{1}(X_t = 1) \neq Y_t) &= \Pr(U_t \leqslant p_1^t \text{ and } U_t > 1/k) + \Pr(U_t > p_1^t \text{ and } U_t \leqslant 1/k) \\
&= \Pr(U_t \in (1/k, p_1^t] \text{ and } p_1^t > 1/k) + \Pr(U_t \in [p_1^t, 1/k] \text{ and } p_1^t \leqslant 1/k) \\
&= \mathbb{E}[|p_1^t - 1/k|].
\end{aligned}$$

Therefore

$$\mathbb{E}[|f(\mathbf{X}) - \mathrm{Bin}(N, 1/k)|] = \mathbb{E}\left[\left|\sum_{t=1}^{N} \mathbb{1}(X_t = 1) - Y_t\right|\right] \leqslant \sum_{t=1}^{N} \mathbb{E}[|p_1^t - 1/k|] \leqslant c\sqrt{n},$$

where this last step used (12). We will now exploit the deviation of $\mathrm{Bin}(N, 1/k)$ from its mean $N/k = n$ to show that $f(\mathbf{X})$ also deviates from n. Since $k \geqslant 2$, we can pick a small enough $\beta > 0$ such that

$$\Pr(\mathrm{Bin}(N, 1/k) \geqslant n + \beta\sqrt{n}) \geqslant \frac{1}{4}.$$

Taking $c = \beta/16$ together with Markov's inequality gives

$$\Pr(\mathrm{Bin}(N, 1/k) - f(\mathbf{X}) \geqslant \beta\sqrt{n}/2) \leqslant \frac{\mathbb{E}[|f(\mathbf{X}) - \mathrm{Bin}(N, 1/k)|]}{\beta\sqrt{n}/2} \leqslant 2c/\beta = \frac{1}{8}.$$

It follows that

$$\begin{aligned}
&\Pr(f(\mathbf{X}) \geqslant n + \beta\sqrt{n}/2) \\
&\qquad \geqslant \Pr(\mathrm{Bin}(N, 1/k) \geqslant n + \beta\sqrt{n} \text{ and } \mathrm{Bin}(N, 1/k) - f(\mathbf{X}) < \beta\sqrt{n}/2) \\
&\qquad \geqslant \Pr(\mathrm{Bin}(N, 1/k) \geqslant n + \beta\sqrt{n}) - \Pr(\mathrm{Bin}(N, 1/k) - f(\mathbf{X}) \geqslant \beta\sqrt{n}/2) \\
&\qquad \geqslant 1/8.
\end{aligned}$$

This concludes the proof of the claim, hence also the proof of this theorem. □

To apply Theorem 11, we need to find a class of digraphs satisfying the necessary spectral conditions. Eulerian tournaments turn out to be such a class, and to establish this we need the following lemma.

Lemma 12. *Every tournament T on an even number of vertices has* $\det(A_T) \neq 0$.

Proof. Suppose T is a tournament on n vertices. We will prove that $\det(A_T)$ is an odd number and hence is not 0. Let I be the $n \times n$ identity matrix and J be the $n \times n$ all 1 matrix. Notice that A_T and $J - I$ are the same over the field $\mathbb{Z}/2\mathbb{Z}$. Because J has eigenvalue n with multiplicity 1 and 0 with multiplicity $n - 1$, we have

$$\det(A_T) \equiv \det(J - I) \equiv (-1)^{n-1}(n - 1) \equiv 1 \mod 2,$$

where this last step holds when n is even. □

Lemma 13. *All Eulerian tournaments D satisfy* $\mathrm{Null}(A_D) = \mathrm{span}(\mathbf{1})$.

Proof. First notice that Eulerian tournaments have an odd number of vertices. Clearly $\mathbf{1} \in \mathrm{Null}(A_D)$. Let T be the tournament obtained by deleting an arbitrary vertex of D. By Lemma 12,

$$\mathrm{rank}(A_D) \geqslant \mathrm{rank}(A_T) = |V(D)| - 1,$$

which means $\dim(\mathrm{Null}(A_D)) = 1$. □

Proof of Theorem 4. The upper bound follows from Theorem 2 while the lower bound follows from Theorem 11 and Lemma 13. □

4 Oblivious Strategies

Inspired by the optimal strategy for semi-restricted RPS, We say that a strategy for Rei $\mathcal{R}$ is *oblivious* if the random variable $\mathcal{R}(r)$ depends only on $\mathrm{supp}(r)$ for all r, i.e. if Rei plays based only on the set of options she can play without taking into account how many times she can perform each option. We say that a digraph D is *oblivious* if Rei has an optimal strategy in the semi-restricted D-game which is oblivious. We emphasize that D being oblivious only guarantees that there exists at least one oblivious optimal strategy, not that every optimal strategy is oblivious.

Naively, oblivious strategies seem like they would be ineffective. However, in Section 2 we saw that every digraph on at most 3 vertices is oblivious, and the unique optimal strategy in the game considered in [11] was also oblivious. Thus at this point one might guess that every digraph is oblivious.

Unfortunately this turns out not to be the case. In fact, we will show that almost every tournament and almost every Eulerian tournament fails to be oblivious. We do this through the following proposition.

Proposition 14. *Let D be a tournament. If there exists a set $S \subseteq V(D)$ with $|S|$ even and with $N^+(v) \cap S \not\subseteq N^+(w) \cap S$ for all $v \in S$, $w \in V(D)$ and $v \neq w$, then D is not oblivious.*

To prove this, we need the following technical result.

Lemma 15. *Let D be an oblivious digraph and $\mathcal{R}$ an oblivious optimal strategy for Rei in the semi-restricted D-game. Let $S \subseteq V(D)$, and let p be the probability vector satisfying $p_u = \Pr(\mathcal{R}(r) = u)$ whenever $\mathrm{supp}(r) = S$. If $v \in S$ is such that there exists a non-negative vector q with $\mathrm{supp}(q) = S$ and $\sum_{u \in N^+(v)} q_u - \sum_{u \in N^-(v)} q_u = \max_w \{\sum_{u \in N^+(w)} q_u - \sum_{u \in N^-(w)} q_u\}$, then we have*

$$\sum_{u \in N^+(v)} p_u - \sum_{u \in N^-(v)} p_u = \max_w \left\{ \sum_{u \in N^+(w)} p_u - \sum_{u \in N^-(w)} p_u \right\}.$$

Let us pause for a moment to discuss what this lemma is saying and where it comes from. Given that Rei plays according to the distribution p when $\mathrm{supp}(r) = S$, Norman will always choose a v achieving $\max_w \left\{\sum_{u\in N^+(w)} p_u - \sum_{u\in N^-(w)} p_u\right\}$. Observe that some vertices v will never achieve this maximum regardless of what p is, e.g. if $N^+(v)$ is disjoint from S and $N^-(v) \neq \emptyset$. Slightly more generally (and almost trivially), the only v which could possibly achieve this maximum for (some) p are those vertices v such that there exists a non-negative vector q with $\mathrm{supp}(q) \subseteq S$ such that v achieves $\max_w \left\{\sum_{u\in N^+(w)} q_u - \sum_{u\in N^-(w)} q_u\right\}$. If the hypothesis of the lemma had the weaker condition $\mathrm{supp}(q) \subseteq S$ rather than $\mathrm{supp}(q) = S$, then the conclusion could be interpreted as saying "any v which could possibly achieve the maximum for p does achieve the maximum", and as such the current lemma can be viewed as a slightly weaker version of this statement.

Proof. Assume v, q are as in the hypothesis of the proposition, and let n be a large integer. Define the vector r by having[3] $r_u = q_u n$ for all u, noting that $\mathrm{supp}(r) = \mathrm{supp}(q) = S$. Let $\gamma = \min_u q_u/p_u$, where we let $q_u/p_u := \infty$ if $p_u = 0$. Note that $0 < \gamma < \infty$ since $\mathrm{supp}(q) = S$ and $p \neq 0$.

Claim 16. *With probability tending towards 1 as n tends towards infinity, an option will be depleted after $(\gamma + o(1))n$ rounds, at which point every vertex u has $(q_u - p_u\gamma)n + o(n)$ actions remaining in expectation.*

Proof of claim. The analysis is similar to the one performed in the proof of Theorem 3.

Rei generates an infinite random string $\pi = \pi_1\pi_2\ldots$ where $\pi_t = u$ independently with probability p_u. In the t-th round Rei will play π_t until an option is depleted. Let T_u denote the smallest integer t such that $|\{s \leqslant t : \pi_s = u\}| = r_u$, and let $T = \min_{u\in\mathrm{supp}(r)} T_u$. Note that T is the number of rounds until an option is depleted.

For any vertex u, we have $T_u \sim X_1 + \ldots + X_{r_u}$ where $X_i \sim \mathrm{Geo}(p_u)$ iid. Therefore $\mathbb{E}T_u = r_u/p_u$ and $\mathrm{Var}(T_u) = o((\mathbb{E}T_u)^2)$. It follows from Chebyshev's inequality that, for example,

$$\Pr(|T_u - \mathbb{E}T_u| \geqslant (\mathbb{E}T_u)^{2/3}) = o(1).$$

By a union bound, we see that with high probability $T_u = (1 + o(1))\mathbb{E}T_u$ for all u, and hence $T = (\gamma + o(1))n$ with high probability.

Moreover, by a routine application of the second moment method and union bound, with high probability the number of times u appears in the first $(\gamma + o(1))n$ letters in π is $(\gamma + o(1))p_u n$ for all u, so after the first vertex is depleted, each vertex u has $(q_u - p_u\gamma)n + o(n)$ actions remaining in expectation. □

Consider the following (possibly non-optimal) strategy for Norman: pick an arbitrary vertex w and play this until some option is depleted, then play v for the rest of the game.

[3]Strictly speaking we should define $r_u = \lceil q_u n\rceil$ to make r a restriction vector, but this distinction will not affect our analysis.

By the claim above, the expected score for Norman under this strategy is

$$\left(\sum_{u\in N^+(w)} p_u - \sum_{u\in N^-(w)} p_u\right)\gamma n + \left(\sum_{u\in N^+(v)} (q_u - p_u\gamma)n - \sum_{u\in N^-(v)} (q_u - p_u\gamma)n\right) + o(n).$$

The quantity above is a lower bound for $S_D(r)$. By using Theorem 3 and the hypothesis on q, we have

$$S_D(r) \leqslant \left(\sum_{u\in N^+(v)} q_u - \sum_{u\in N^-(v)} q_u\right) n + O(n^{2/3}).$$

Comparing these two inequalities for $S_D(r)$ gives

$$\left(\sum_{u\in N^+(w)} p_u - \sum_{u\in N^-(w)} p_u\right)\gamma n \leqslant \left(\sum_{u\in N^+(v)} p_u - \sum_{u\in N^-(v)} p_u\right)\gamma n + o(n).$$

This is only possible if $\sum_{u\in N^+(w)} p_u - \sum_{u\in N^-(w)} p_u \leqslant \sum_{u\in N^+(v)} p_u - \sum_{u\in N^-(v)} p_u$. As w was arbitrary, this implies the result. □

We can now prove Proposition 14.

Proof of Proposition 14. Assume for contradiction that an oblivious optimal strategy $\mathcal{R}$ existed, and let p be the probability vector satisfying $p_u = \Pr(\mathcal{R}(r) = u)$ whenever $\operatorname{supp}(r) = S$.

For each $v \in S$, consider the non-negative vector q with $q_u = 2|V(D)|$ for $u \in N^+(v) \cap S$, $q_u = 1$ for $u \in S \setminus N^+(v)$, and $q_u = 0$ otherwise. By hypothesis, for all $w \neq v$ we have

$$|N^+(w) \cap (N^+(v) \cap S)| \leqslant |N^+(v) \cap S| - 1.$$

This implies for all $w \neq v$ that

$$\begin{aligned}\sum_{u\in N^+(w)} q_u - \sum_{u\in N^-(w)} q_u \leqslant \sum_{u\in N^+(w)} q_u &\leqslant |N^+(w) \cap (N^+(v) \cap S)| \cdot 2|V(D)| + |V(D)| \cdot 1 \\ &\leqslant |N^+(v) \cap S| \cdot 2|V(D)| - |V(D)| \leqslant \sum_{u\in N^+(v)} q_u - \sum_{u\in N^-(v)} q_u.\end{aligned}$$

Because $\operatorname{supp}(q) = S$, Lemma 15 implies

$$\sum_{u\in N^+(v)} p_u - \sum_{u\in N^-(v)} p_u = \max_w \left\{\sum_{u\in N^+(w)} p_u - \sum_{u\in N^-(w)} p_u\right\} \tag{13}$$

for all $v \in S$.

Since $\operatorname{supp}(p) \subseteq S$, we have

$$\sum_{v\in S}\left(\sum_{u\in N^+(v)} p_u - \sum_{u\in N^-(v)} p_u\right) = \sum_{v\in S}\left(\sum_{u\in N^+(v)\cap S} p_u - \sum_{u\in N^-(v)\cap S} p_u\right) = 0.$$

This together with the fact that (13) holds for all $v \in S$ implies that we must have

$$\sum_{u \in N^+(v)} p_u - \sum_{u \in N^-(v)} p_u = \sum_{u \in N^+(v) \cap S} p_u - \sum_{u \in N^-(v) \cap S} p_u = 0 \tag{14}$$

for all $v \in S$. Define $A_D[S]$ to be the submatrix of A_D restricted to the rows and columns indexed by S. Then, (14) states that p restricted to the indices of S is a nullvector for $A_D[S]$, and p being a probability vector means that it is a non-zero nullvector. But from Lemma 12 we know that such a vector p does not exist, giving the desired contradiction. □

Using Proposition 14, we can quickly establish that an exponentially small proportion of tournaments are oblivious.

Theorem 17. *Let $\mathcal{T}_n$ denote the set of tournaments on $[n]$. There exists a constant $c > 0$ such that if n is sufficiently large and D is chosen uniformly at random from $\mathcal{T}_n$, then*

$$\Pr(D \textit{ is oblivious}) \leqslant e^{-cn}.$$

Proof. Let $m = n$ if n is even, and $m = n - 1$ otherwise. Fix any subset $S \subseteq [n]$ of size m. Let B denote the event that there exists $v \in S$, $w \in V(D)$ with $v \neq w$ such that $N^+(v) \cap S \subseteq N^+(w) \cap S$. By Proposition 14, if D is oblivious, then B must occur. Thus it suffices to upper bound $\Pr(B)$.

For any $v \neq w \in V(D)$, it is straightforward to show that

$$\Pr(N^+(v) \cap S \subseteq N^+(w) \cap S) \leqslant (3/4)^{n-3}.$$

Thus by a union bound, we find

$$\Pr(D \text{ is oblivious}) \leqslant \Pr(B) \leqslant n^2 \cdot (3/4)^{n-3},$$

giving the desired result. □

Similarly, almost every Eulerian tournament is not oblivious.

Theorem 18. *Let $\mathcal{E}_n$ denote the set of Eulerian tournaments on $[2n+1]$. If D is chosen uniformly at random from $\mathcal{E}_n$, then*

$$\Pr(D \textit{ is oblivious}) = O(n^{-2}).$$

Proving Theorem 18 requires a significantly more complicated argument than that of Theorem 17. Philosophically, this is because it is easy to sample a uniformly random tournament (one can just choose the orientation of each arc uniformly and independently), but it is much harder to generate an Eulerian tournament uniformly at random, see for example [10] which implicitly provides an efficient probabilistic sampling algorithm using random walks. We get around this issue by invoking a "switching" type argument, the full details of which can be found in Appendix B.

5 Open Problems

Many open questions about semi-restricted games remain. We discuss a few of these in the following subsections.

5.1 Optimal Strategies

In this paper, we determined all of the optimal strategies for semi-restricted D-games when D has at most 3 vertices. It would be of interest to do this for some (non-trivial) infinite family of digraphs as well. Perhaps the simplest such family is the following.

Problem 19. Determine all of the optimal strategies for the semi-restricted P_n-game, where P_n denotes the directed path on n vertices.

In Theorem 18, we showed that almost every Eulerian tournament D fails to have a "simple" (i.e. oblivious) optimal strategy for Rei in the semi-restricted D-game. A closer inspection of the proof shows that these tournaments fail to have such an optimal strategy when Rei has $|V(D)|-1$ remaining options. Thus it is possible (though seemingly unlikely) that Rei has a simple optimal strategy when every option is still available to her.

Problem 20. Is it true that for every Eulerian tournament D, there exists an optimal strategy for Rei in the semi-restricted D-game such that under this strategy, if Rei has restriction vector r with $\mathrm{supp}(r) = V(D)$, then Rei plays each option with probability $1/|V(D)|$?

5.2 Improved Bounds

Theorem 1 shows that the expected score for Norman in semi-restricted RPS is $\Theta(\sqrt{n})$ when both players play optimally. It would be interesting to get a more precise estimate.

Question 21. Does there exist a $c > 0$ such that $S_D(n,n,n) \sim c\sqrt{n}$ when D is the directed 3-cycle? If so, what is c?

Because Theorem 1 gives the optimal strategy for Rei in this game, Question 21 is equivalent to asymptotically determining the expected value of some strange (but explicit) random variable. In particular, it suffices to understand the times T_1, T_2 at which point Rei runs out of 1 or 2 options. It is plausible that these random variables (up to rescaling) are exponentials as in the original birthday and coupon collector problems. We have computed the exact value of $S_D(n,n,n)$ for $n \leqslant 100$, and this data suggests that we may have $c \approx 1.46$.

We proved effective general upper bounds on $S_D(r)$ through Theorems 2 and 3. It is natural to ask if these results can be improved. For example, it might be possible to improve the $O_D(M^{2/3})$ error term in Theorem 3 to $O_D(M^{1/2})$.

Question 22. Is it true that for every digraph D and restriction vector r with $M = \max_v r_v$, we have

$$S_D(r) \leqslant \max_v \left\{ \sum_{u \in N^+(v)} r_u - \sum_{u \in N^-(v)} r_u \right\} + O_D(M^{1/2}).$$

Such a result would be an (optimal) strengthening of Theorem 3, as well as a significant generalization of Theorem 2. The central obstacle with this question is that we do not know what strategy Rei should use so that Norman can obtain at most this many points in expectation.

When D is an Eulerian digraph, Theorem 2 shows that $0 \leqslant S_D(n \cdot \mathbf{1}) \leqslant O_D(\sqrt{n})$. We think that this upper bound might be tight in general.

Question 23. Does every Eulerian digraph with at least one arc satisfy $S_D(n \cdot \mathbf{1}) = \Omega_D(\sqrt{n})$?

A positive answer to this question together with Theorem 2 would show $S_D(n \cdot \mathbf{1}) = \Theta_D(\sqrt{n})$ for all such D.

We can show Question 23 has a positive answer whenever the spectral conditions of Theorem 11 are satisfied. Motivated by this, we ask when exactly these spectral conditions occur. We emphasize that this is a purely linear algebraic question and does not involve any game theory.

Question 24. Which Eulerian digraphs D are such that their skew-adjacency matrix A_D has a nullspace of dimension 1?

For example, Lemma 13 shows that Eulerian tournaments have this property, and we can also show that this holds for certain powers of directed Hamiltonian cycles. Similarly it is not difficult to show that D will fail to satisfy this property if its underlying graph is bipartite, has an even number of vertices, or if there exist distinct vertices v, v' with $N^+(v) = N^+(v')$ and $N^-(v) = N^-(v')$.

We note that the matrix A_D, as well as the essentially equivalent Hermitian matrix $i \cdot A_D$, have been well studied in the literature, see for example [5, 8], and it is possible that the results and techniques used in these papers could give insight into the answer to Question 24.

Acknowledgements

The authors are indebted to Yuanfan Wang for many fruitful discussions on this topic, as well as the referees for their careful reading and helpful suggestions. We would also like to thank Jack J. Garzella, as well as the Mathematics Graduate Student Council at UCSD, for organizing the workshop where this work was initiated. We thank Krystal Guo and Marcus Michelen for useful discussions regarding the skew-adjacency matrix A_D, Lutz Warnke for discussions regarding switching arguments, and Joel Spencer for discussions about Question 21.

The first author was supported by the National Science Foundation Graduate Research Fellowship Grant DGE-1650112. The second author was supported by the National Science Foundation Grant DMS-2225631. The third author was supported by the National Science Foundation Grant DMS-1800746.

References

[1] Y. Alimohammadi, P. Diaconis, M. Roghani, and A. Saberi. Sequential importance sampling for estimating expectations over the space of perfect matchings. *Annals of Applied Probability*, to appear.

[2] P. Diaconis and R. Graham. The analysis of sequential experiments with feedback to subjects. *Annals of Statistics*, 9(1):3–23, 1981.

[3] P. Diaconis, R. Graham, X. He, and S. Spiro. Card guessing with partial feedback. *Combinatorics, Probability and Computing*, 31(1):1–20, 2022.

[4] N. Fukumoto. *Tobaku Mokushiroku: Kaiji*, volume 1. Kodansha Ltd., 1996.

[5] K. Guo and B. Mohar. Hermitian adjacency matrix of digraphs and mixed graphs. *Journal of Graph Theory*, 85(1):217–248, 2017.

[6] J. He and A. Ottolini. Card guessing and the birthday problem for sampling without replacement. arXiv:2108.07355, 2021.

[7] T. Krityakierne and T. A. Thanatipanonda. No feedback? no worries! the art of guessing the right card. arXiv:2205.08793, 2022.

[8] J. Liu and X. Li. Hermitian-adjacency matrices and Hermitian energies of mixed graphs. *Linear Algebra and its Applications*, 466:182–207, 2015.

[9] P. Liu. On card guessing game with one time riffle shuffle and complete feedback. *Discrete Applied Mathematics*, 288:270–278, 2021.

[10] B. D. McKay and R. W. Robinson. Asymptotic enumeration of Eulerian circuits in the complete graph. *Combinatorics, Probability and Computing*, 7(4):437–449, 1998.

[11] S. Spiro. Online card games. *Electronic Journal of Probability*, 27:1–15, 2022.

[12] J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 60th anniversary edition, 2007.

A Proof of Theorem 8

This section is dedicated to proving Theorem 8, which we recall says that if D is the directed 3-path $1 \to 2 \to 3$, then the optimal strategies for Rei in the semi-restricted D-game are exactly those $\mathcal{R}$ satisfying $\Pr(\mathcal{R}(r) = 3) = 1/2$ whenever $\{3\} \subsetneq \operatorname{supp}(r)$.

Proof of Theorem 8. Let $\mathcal{R}$ be an arbitrary optimal strategy for Rei, and as before we argue by induction that $\Pr(\mathcal{R}(r) = 3) = 1/2$ whenever $\{3\} \subsetneq \operatorname{supp}(r)$. We fix an arbitrary r with $\{3\} \subsetneq \operatorname{supp}(r)$ and write $\Pr(\mathcal{R}(r) = u) = p_u$ for each $u \in \{1, 2, 3\}$.

If $p_2 = 0$, then by applying the techniques in the proof of Theorem 7, we can argue $p_1 = p_3 = 1/2$ as desired. Similarly if $p_1 = 0$, we can argue $p_2 = p_3 = 1/2$. As such we skip these arguments and consider only the case $p_1, p_2 > 0$. Our aim is to show $p_2 = p_3 - p_1$, which implies $p_3 = 1/2$ as desired.

If $p_2 > p_3 - p_1$, let

$$\epsilon = \min\{p_2, p_2 - (p_3 - p_1)\} > 0$$

and consider a different strategy $\mathcal{R}'$ for Rei with

$$\Pr(\mathcal{R}'(r) = 1) = p_1 + \frac{\epsilon}{2},\ \Pr(\mathcal{R}'(r) = 2) = p_2 - \epsilon,\ \Pr(\mathcal{R}'(r) = 3) = p_3 + \frac{\epsilon}{2},$$

and $\Pr(\mathcal{R}'(r') = u) = \Pr(\mathcal{R}(r') = u)$ for all vertices u and $r' \neq r$. Note that, by our choice of ϵ, we have

$$\max\left\{p_2 - \epsilon, \left(p_1 + \frac{\epsilon}{2}\right) - \left(p_3 + \frac{\epsilon}{2}\right), -(p_2 - \epsilon)\right\} = \max\{p_2, p_3 - p_1, -p_2\} - \epsilon.$$

Using Fact 5 and that $\mathcal{R}$ is an optimal strategy, we find

$$S_D(r; \mathcal{R}') - S_D(r; \mathcal{R}) = \epsilon\left(\frac{1}{2}S_D(r - \delta_1) + \frac{1}{2}S_D(r - \delta_3) - S_D(r - \delta_2) - 1\right).$$

By Lemma 6 and the structure of D, we have

$$S_D(r - \delta_1) \leqslant S_D(r - \delta_2) + 1 \text{ and } S_D(r - \delta_3) < S_D(r - \delta_2) + 1,$$

which implies $S_D(r; \mathcal{R}') - S_D(r; \mathcal{R}) < 0$. This contradicts $\mathcal{R}$ being optimal, so we can not have $p_2 > p_3 - p_1$. If $p_2 < p_3 - p_1$, we can take $\epsilon = (p_3 - p_1) - p_2 > 0$ and consider another strategy $\mathcal{R}'$ for Rei with

$$\Pr(\mathcal{R}'(r) = 1) = p_1,\ \Pr(\mathcal{R}'(r) = 2) = p_2 + \frac{\epsilon}{2},\ \Pr(\mathcal{R}'(r) = 3) = p_3 - \frac{\epsilon}{2},$$

with $\mathcal{R}'$ coinciding with $\mathcal{R}$ on every other vector r'. Similarly, $\mathcal{R}'$ will contradict the fact that $\mathcal{R}$ is optimal, so we must have $p_2 = p_3 - p_1$ as wanted. This proves that every optimal strategy $\mathcal{R}$ satisfies $\Pr(\mathcal{R}(r) = 3) = 1/2$ when $\{3\} \subsetneq \operatorname{supp}(r)$. It remains to show that all such strategies are optimal, i.e. that $S_D(r; \mathcal{R}) = S_D(r; \mathcal{R}')$ for any strategies $\mathcal{R}, \mathcal{R}'$ of this form.

Take $\mathcal{R}$ to be the unique strategy for Rei defined by $\Pr(\mathcal{R}(r) = 3) = 1/2$ whenever $\{3\} \subsetneq \operatorname{supp}(r)$ and $\Pr(\mathcal{R}(r) = 2) = 0$ whenever $1 \in \operatorname{supp}(r)$. We shall prove by induction on $r_1 + r_2 + r_3$ that

- $S_D(r;\mathcal{R}) = S_D(r;\mathcal{R}')$ for an arbitrary strategy $\mathcal{R}'$ satisfying $\Pr(\mathcal{R}'(r) = 3) = 1/2$ whenever $\{3\} \subsetneq \mathrm{supp}(r)$;
- $S_D(r) = S_D(r - \delta_1 + \delta_2) - 1$ whenever $r_1 > 0$.

The base case when $r_1+r_2+r_3 = 0$ is trivial. Suppose we have proved these two identities for all r' with $r'_1 + r'_2 + r'_3 < r_1 + r_2 + r_3$.

First we argue $S_D(r;\mathcal{R}) = S_D(r;\mathcal{R}')$. Let us write $\Pr(\mathcal{R}(r) = u) = p_u$ and $\Pr(\mathcal{R}'(r) = u) = p'_u$ for all u. If either $r_1 = 0$ or $r_2 = 0$, by the description of $\mathcal{R}$ and $\mathcal{R}'$, we know they coincide on all vectors r' with $r'_1 \leqslant r_1$, $r'_2 \leqslant r_2$, and $r'_3 \leqslant r_3$. So we clearly have the desired identity in this case. If $r_1, r_2 > 0$ and $r_3 = 0$, by the description of $\mathcal{R}, \mathcal{R}'$, we have $p_1 = 1$, $p_2 = p_3 = p'_3 = 0$, and $p'_1 = 1 - p'_2$. Using Fact 5 we can compute

$$\begin{aligned} S_D(r;\mathcal{R}) - S_D(r;\mathcal{R}') &= S_D(r - \delta_1) - (p'_2 + (1 - p'_2)S_D(r - \delta_1) + p'_2 S_D(r - \delta_2)) \\ &= -p'_2(1 - S_D(r - \delta_1) + S_D(r - \delta_2)) = 0, \end{aligned}$$

where the last equality follows from the second identity from the inductive hypothesis. If $r_1, r_2, r_3 > 0$, by the description of $\mathcal{R}, \mathcal{R}'$, we have $p_1 = 1/2$, $p_2 = 0$, $p_3 = p'_3 = 1/2$, and $p'_1 = 1/2 - p'_2$. Again, by Fact 5 we can compute

$$S_D(r;\mathcal{R}) - S_D(r;\mathcal{R}') = -p'_2(1 - S_D(r - \delta_1) + S_D(r - \delta_2)) = 0.$$

So we conclude the first wanted identity for our inductive process.

Next we argue $S_D(r) = S_D(r - \delta_1 + \delta_2) - 1$ under the condition $r_1 > 0$. Note that we have already showed $S_D(r;\mathcal{R}) = S_D(r;\mathcal{R}')$, and similarly we have $S_D(r - \delta_1 + \delta_2;\mathcal{R}) = S_D(r - \delta_1 + \delta_2;\mathcal{R}')$. So by the first half of this proof, we have

$$S_D(r) = S_D(r;\mathcal{R}) \quad \text{and} \quad S_D(r - \delta_1 + \delta_2) = S_D(r - \delta_1 + \delta_2;\mathcal{R}).$$

Now, if $r_3 = 0$, it is easy to check $S_D(r) = r_2 = S_D(r - \delta_1 + \delta_2) - 1$ as wanted. If $r_3 > 0$ and $r_1 = 1$, by Fact 5 and the description of $\mathcal{R}$, we have

$$\begin{aligned} S_D(r) &= \frac{1}{2}S_D(r - \delta_1) + \frac{1}{2}S_D(r - \delta_3), \\ S_D(r - \delta_1 + \delta_2) &= \frac{1}{2} + \frac{1}{2}S_D(r - \delta_1) + \frac{1}{2}S_D(r - \delta_1 + \delta_2 - \delta_3). \end{aligned}$$

Together with the inductive hypothesis $S_D(r - \delta_3) = S_D(r - \delta_1 + \delta_2 - \delta_3) - 1$, we can conclude the wanted identity. If $r_3 > 0$ and $r_1 > 1$, by Fact 5 and the description of $\mathcal{R}$, we have

$$\begin{aligned} S_D(r) &= \frac{1}{2}S_D(r - \delta_1) + \frac{1}{2}S_D(r - \delta_3), \\ S_D(r - \delta_1 + \delta_2) &= \frac{1}{2}S_D(r - 2\delta_1 + \delta_2) + \frac{1}{2}S_D(r - \delta_1 + \delta_2 - \delta_3). \end{aligned}$$

Again, applying the inductive hypothesis for $S_D(r - \delta_1)$ and $S_D(r - \delta_3)$ respectively, we can conclude $S_D(r) = r_2 = S_D(r - \delta_1 + \delta_2) - 1$ as wanted. Hence we conclude the inductive process, which completes the proof. □

B Proof of Theorem 18

This section is dedicated to proving Theorem 18, which we recall says that if $\mathcal{E}_n$ is the set of Eulerian tournaments on $[2n+1]$ and D is chosen uniformly at random from $\mathcal{E}_n$, then

$$\Pr(D \text{ is oblivious}) = O(n^{-2}).$$

The main lemma we need is the following.

Lemma 25. *Let $\mathcal{E}_n$ denote the set of Eulerian tournaments on $[2n+1]$, and for $v, w \in [2n+1]$, let $\mathcal{E}_n^{v,w} \subseteq \mathcal{E}_n$ denote the set of Eulerian tournaments D with $vw \in E(D)$ and with $|N^+(v) \cap N^+(w)| = n-1$. If $n \geqslant 16$, then*

$$|\mathcal{E}_n^{v,w}| \leqslant 2^{14} n^{-4} |\mathcal{E}_n|.$$

Proof. It is not difficult to see that for each $D \in \mathcal{E}_n^{v,w}$, there exists a vertex u_D and disjoint sets A_D, B_D of size $n-1$ such that $N^+(v) \cap N^+(w) = A_D$, $N^-(v) \cap N^-(w) = B_D$, and $v \to w \to u_D \to v$. We use this structural result together with a "switching" type argument to prove the bound.

Given $D \in \mathcal{E}_n^{v,w}$, we say that a pair of arcs $\{a_1b_1, a_2b_2\}$ is *D-valid* if these are vertex disjoint arcs in D with $a_1, a_2 \in A_D$ and $b_1, b_2 \in B_D$. Given a D-valid pair $\{a_1b_1, a_2b_2\}$, we define $D[a_1b_1, a_2b_2]$ to be the digraph which has the arcs $b_1 \to a_1 \to v \to b_1$ and $b_2 \to a_2 \to v \to b_2$ and which otherwise agrees with D. It is not difficult to see that $D[a_1b_1, a_2b_2]$ is an Eulerian tournament because D was Eulerian and $\{a_1b_1, a_2b_2\}$ was D-valid.

With the above definition in mind, we construct an auxiliary bipartite graph G with vertex set $\mathcal{E}_n^{v,w} \sqcup \mathcal{E}_n$ by having $D \in \mathcal{E}_n^{v,w}$ and $D' \in \mathcal{E}_n$ adjacent in G if and only if $D' = D[a_1b_1, a_2b_2]$ for some D-valid pair $\{a_1b_1, a_2b_2\}$.

Claim 26. *Each $D \in \mathcal{E}_n^{v,w}$ has $\deg_G(D) \geqslant 2^{-7} n^4$.*

Proof. Since $D[a_1b_1, a_2b_2]$ is a distinct digraph for each distinct D-valid pair $\{a_1b_1, a_2b_2\}$, the claim is equivalent to saying each $D \in \mathcal{E}_n^{v,w}$ has at least $2^{-7}n^2$ many D-valid pairs $\{a_1b_1, a_2b_2\}$.

Given two (possibly non-disjoint) sets $S, T \subseteq V(D)$, let $e(S,T)$ be the number of arcs $s \to t$ in D with $s \in S$ and $t \in T$. Because D is Eulerian and $|A_D| = n-1$, we have

$$\begin{aligned} n(n-1) = e(A_D, V(D)) &= e(A_D, A_D) + e(A_D, \{v,w\}) + e(A_D, \{u_D\}) + e(A_D, B_D) \\ &\leqslant \binom{n-1}{2} + 0 + (n-1) + e(A_D, B_D), \end{aligned}$$

which implies $e(A_D, B_D) \geqslant \frac{1}{2}(n-1)(n-4)$. Thus there exist at least $\frac{1}{2}(n-1)(n-4)$ choices for arcs a_1b_1 from A_D to B_D, and given such an arc, there exist at least

$$e(A_D, B_D) - (n-1) \cdot 2 \geqslant \frac{1}{2}(n-1)(n-8)$$

arcs a_2b_2 with $a_1 \neq a_2, b_1 \neq b_2$. In total then the number of ways we can construct an (ordered) D-valid pair (a_1b_1, a_2b_2) is at least

$$\frac{1}{2}(n-1)(n-4) \cdot \frac{1}{2}(n-1)(n-8) \geqslant 2^{-6}n^4$$

for $n \geqslant 16$. This double counts the number of (unordered) D-valid pairs $\{a_1b_1, a_2b_2\}$, so dividing this quantity by 2 gives the desired lower bound. □

Claim 27. *Each $D' \in \mathcal{E}_n$ has $\deg_G(D') \leqslant 120$.*

Proof. Observe that if $D \in \mathcal{E}_n^{v,w}$ and $\{a_1b_1, a_2b_2\}$ is D-valid, then there are exactly 5 vertices $x \in V(D[a_1b_1, a_2b_2])$ which have one arc to $\{v, w\}$ and one arc from $\{v, w\}$, namely this holds for a_1, b_1, a_2, b_2, u_D. Thus $\deg_G(D') = 0$ if $D' \in \mathcal{E}_n$ does not have 5 vertices with this property, and otherwise there are trivially at most $5! = 120$ digraphs D which could have $D' = D[a_1b_1, a_2b_2]$ (namely by choosing which of its 5 vertices to play the roles of a, a', b, b', u_D). □

With these two claims, we have

$$2^{-7}n^2|\mathcal{E}_n^{v,w}| \leqslant \sum_{D \in \mathcal{E}_n^{v,w}} \deg_G(D) = \sum_{D' \in \mathcal{E}_n} \deg_G(D') \leqslant 120|\mathcal{E}_n|,$$

and rearranging this inequality gives the result. □

We can now prove almost every Eulerian tournament is not oblivious.

Proof of Theorem 18. Let $\mathcal{E}_n^* \subseteq \mathcal{E}_n$ be the set of oblivious Eulerian digraphs on $[2n+1]$. Then the theorem statement is equivalent to saying

$$|\mathcal{E}_n^*| = O(n^{-2}|\mathcal{E}_n|).$$

Fix an arbitrary set $S \subseteq [2n+1]$ of $2n$ vertices, and let $\mathcal{E}_n^{**}$ be the set of digraphs which have $N^+(v) \cap S \subseteq N^+(w) \cap S$ for some distinct $v, w \in [2n+1]$. We claim that

$$\mathcal{E}_n^* \subseteq \mathcal{E}_n^{**} \subseteq \bigcup_{v,w} \mathcal{E}_n^{v,w}.$$

Indeed, $\mathcal{E}_n^* \subseteq \mathcal{E}_n^{**}$ follows from Proposition 14. Let $D \in \mathcal{E}_n^{**}$, say with $N^+(v) \cap S \subseteq N^+(w) \cap S$. Note that $|N^+(v) \cap S| \geqslant |N^+(v)| - 1 = n - 1$, so in particular $|N^+(v) \cap N^+(w)| \geqslant n - 1$. We can not have $N^+(v) = N^+(w)$ since either $vw \in E(D)$ or $wv \in E(D)$, so we must have $|N^+(v) \cap N^+(w)| = n - 1$. Thus $D \in \mathcal{E}_n^{v,w} \cup \mathcal{E}_n^{w,v}$, proving the claim.

With this claim, we have

$$|\mathcal{E}_n^*| \leqslant |\mathcal{E}_n^{**}| \leqslant \left|\bigcup_{v,w} \mathcal{E}_n^{v,w}\right| = O(n^{-2}|\mathcal{E}_n|),$$

where this last step used Lemma 25. We conclude the result. □

We note that it is likely that a more sophisticated argument could be used to improve the bound of Theorem 18 to show that it is exponentially unlikely for a random Eulerian tournament to be oblivious.

C Restricted Games

This section is devoted to studying games where both players are restricted. In particular, we look at the case when two players Alice and Bob are given restriction vectors a, b with $\sum a_i = \sum b_j = N$ and play some zero sum game a total of N times, with e.g. Alice being forced to use each option i a total of a_i times. We begin with some formal definitions analogous to the definitions of Section 2.

We use G to denote a simultaneous zero-sum game played by Alice and Bob where O_A (O_B resp.) consists of the set of options Alice (Bob resp.) can play at each round. We let $G(i, j)$ denote the score that Alice receives in this game if Alice plays i and Bob plays j. We say that (a, b) is a pair of restriction vectors if a, b are non-negative integral vectors indexed by O_A, O_B with $\sum a_i = \sum b_j$.

Given a game G as above, a *strategy for Alice in the restricted G-game* is a function $\mathcal{A}$ from pairs of restriction vectors (a, b) to random variables such that $\text{supp}(\mathcal{A}(a, b)) \subseteq \text{supp}(a)$ for all (a, b). A *strategy for Bob* $\mathcal{B}$ is defined analogously by requiring that $\text{supp}(\mathcal{B}(a, b)) \subseteq \text{supp}(b)$. Recall that δ_i is a restriction vector with value 1 on entry i and 0 elsewhere. Given a pair of strategies $\mathcal{A}, \mathcal{B}$ for Alice and Bob and restriction vectors (a, b), we define the *score* (for Alice) $\mathbf{S}_G(a, b; \mathcal{A}, \mathcal{B})$ by $\mathbf{S}_G(a, b; \mathcal{A}, \mathcal{B}) = 0$ if $\sum a_i = \sum b_j = 0$, and otherwise

$$\mathbf{S}_G(a, b; \mathcal{A}, \mathcal{B}) = G(\mathcal{A}(a, b), \mathcal{B}(a, b)) + \mathbf{S}_G(a - \delta_{\mathcal{A}(a,b)}, b - \delta_{\mathcal{B}(a,b)}; \mathcal{A}, \mathcal{B}).$$

We define the *expected score* $S_G(a, b; \mathcal{A}, \mathcal{B}) = \mathbb{E}[\mathbf{S}_G(a, b; \mathcal{A}, \mathcal{B})]$. We have the following analogue of Fact 5 for restricted games.

Fact 28. *Let $\mathcal{A}$ and $\mathcal{B}$ be strategies for Alice and Bob respectively. If (a, b) is a restriction vector pair with $p_i = \Pr(\mathcal{A}(a, b) = i)$ and $q_j = \Pr(\mathcal{B}(a, b) = j)$ for each $i \in O_A, j \in O_B$, then*

$$S_G(a, b; \mathcal{A}, \mathcal{B}) = \sum_i \sum_j p_i q_j (G(i, j) + S_G(a - \delta_i, b - \delta_j; \mathcal{A}, \mathcal{B})).$$

A strategy for Alice $\mathcal{A}$ is *optimal* if for every strategy $\mathcal{B}$ for Bob we have $S_G(a, b; \mathcal{A}, \mathcal{B}) = \max_{\mathcal{A}'} S_G(a, b; \mathcal{A}', \mathcal{B})$, and we similarly define what it means for $\mathcal{B}$ to be an optimal strategy for Bob. We define the *uniform strategy for Alice* $\mathcal{U}_A$ by having $\Pr(\mathcal{U}_A(a, b) = i) = \frac{a_i}{\sum_{i'} a_{i'}}$ and we similarly define the uniform strategy for Bob $\mathcal{U}_B$.

Proposition 29. *In the restricted G-game, $\mathcal{U}_A$ is an optimal strategy for Alice and $\mathcal{U}_B$ is an optimal strategy for Bob.*

Proof. We will prove by induction on N that if a, b are restriction vectors with $\sum_i a_i = \sum_j b_j = N$ and if $\mathcal{A}$ is any strategy for Alice, then

$$S_G(a, b; \mathcal{A}, \mathcal{U}_B) = \sum_{i,j} \frac{G(i, j) a_i b_j}{N} =: M(a, b).$$

The base case when $N = 1$ is obvious. For the inductive process, suppose there exists $i \in O_A$ such that $\Pr(\mathcal{A}(a,b) = i) = 1$. By Fact 28, we have

$$S_G(a,b;\mathcal{A},\mathcal{U}_B) = \sum_j \frac{b_j}{N} \cdot (G(i,j) + S_G(a-\delta_i, b-\delta_j; \mathcal{A},\mathcal{U}_B)). \tag{15}$$

By the inductive hypothesis, we have

$$\begin{aligned} S_G(a-\delta_i, b-\delta_j; \mathcal{A},\mathcal{U}_B) &= \frac{1}{N-1}\left(\sum_{i'\neq i, j'\neq j} G(i',j')a_{i'}b_{j'} + \sum_{i'\neq i} G(i',j)a_{i'}(b_j-1) \right. \\ &\qquad \left. + \sum_{j'\neq j} G(i,j')(a_i-1)b_{j'} + G(i,j)(a_i-1)(b_j-1) \right) \\ &= \frac{1}{N-1}\left(\sum_{i',j'} G(i',j')a_{i'}b_{j'} - \sum_{i'} G(i',j)a_{i'} - \sum_{j'} G(i,j')b_{j'} + G(i,j) \right). \end{aligned}$$

Plugging this into (15) and using $\sum_j cb_j = cN$ for any constant c, we have

$$\begin{aligned} N(N-1)\cdot S_G(a,b;\mathcal{A},\mathcal{U}_B) &= \sum_j b_j \left((N-1)G(i,j) + \sum_{i',j'} G(i',j')a_{i'}b_{j'} \right. \\ &\qquad \left. - \sum_{i'} G(i',j)a_{i'} - \sum_{j'} G(i,j')b_{j'} + G(i,j) \right) \\ &= (N-1)\sum_{i',j'} G(i',j')a_{i'}b_{j'}. \end{aligned}$$

Dividing both sides by $N(N-1)$ gives $S_G(a,b;\mathcal{A},\mathcal{U}_B) = M(a,b)$ under our assumption $\Pr(\mathcal{A}(a,b) = i) = 1$. Since this holds regardless of which deterministic option i Alice plays, we see that $S_G(a,b;\mathcal{A},\mathcal{U}) = M(a,b)$ for any (mixed) strategy $\mathcal{A}$ for Alice, concluding the inductive proof.

As we have shown, Bob can use strategy $\mathcal{U}$ to guarantee that Alice's expected score is at most $M(a,b)$. A symmetric argument shows that Alice can guarantee at least $M(a,b)$ points by using her uniform strategy. Thus $M(a,b)$ is the expected score if both players play optimally, and in particular $\mathcal{U}_A$ and $\mathcal{U}_B$ are optimal strategies for Alice and Bob respectively. □

We note that in general $\mathcal{U}_A$ and $\mathcal{U}_B$ will not be the unique optimal strategies for the restricted G-game. Indeed, if $G(i,j) = 0$ for all i,j then every strategy is optimal, and more generally there will be multiple optimal strategies if there exist $i \neq i'$ such that $G(i,j) = G(i',j)$ for all j.