

Achieving the Highest Possible Elo Rating

Rikhav Shah

February 2024

Abstract

Elo rating systems measure the approximate skill of each competitor in a game or sport. A competitor's rating increases when they win and decreases when they lose. Increasing one's rating can be difficult work; one must hone their skills and consistently beat the competition. Alternatively, with enough money you can rig the outcome of games to boost your rating. This paper poses a natural question for Elo rating systems: say you manage to get together n people (including yourself) and acquire enough money to rig k games. How high can you get your rating, asymptotically in k ? In this setting, the people you gathered aren't very interested in the game, and will only play if you pay them to. This paper resolves the question for $n = 2$ up to constant additive error, and provide close upper and lower bounds for all other n , including for n growing arbitrarily with k . There is phase transition at $n = k^{1/3}$: there is a huge increase in the highest possible Elo rating from $n = 2$ to $n = k^{1/3}$, but (depending on the particular Elo system used) little-to-no increase for any higher n . Past the transition point $n > k^{1/3}$, the highest possible Elo is at least $\Theta(k^{1/3})$. The corresponding upper bound depends on the particular system used, but for the standard Elo system, is $\Theta(k^{1/3} \log(k)^{1/3})$.

1 Introduction

The Elo Rating system was proposed by Arpad Elo in the mid 20th century to estimate the relative skill of chess players [2]. It was quickly adopted by the international chess community, and in the decades since has seen adoption in many competitive contexts. This paper considers a simple combinatorial question about the Elo system. To the author's knowledge, this is the first time this question has been posed in print: **given n players starting with equal rating, what is the highest a player could be rated after a total of k games are played?**

We begin with a definition of the system then provide its motivation. Each player is given some 'rating' value (measured in 'points' or simply 'Elo'), which updates as they play games. These rating points are somewhat analogous to poker chips: when player A and player B play a game, they each place some of their rating points into a pot. In the case of a draw, the players split the pot evenly. If one player wins, they take the entire pot. The heart of the Elo system is dictating how many points each player must ante up. To do so, each implementation of the system specifies a 'pot function' σ satisfying

1. σ is non-negative and monotonically increasing, and
 2. $\sigma(z) + \sigma(-z) = 1$ for all $z \in \mathbb{R}$.
- (1)

Let r_A and r_B be the ratings of players A and B respectively. When players A and B play, the number of points they ante up are $K \cdot \sigma(r_A - r_B)$ and $K \cdot \sigma(r_B - r_A)$ respectively, for a total pot size of K . Players are allowed to go into debt if they don't have the required points, i.e. negative ratings are perfectly fine. The value of K itself is a parameter of the system. The resulting rating updates for different outcomes of a game

between A and B are as follows:

	A wins	A and B draw	B wins
$r'_A =$	$r_A + K \cdot \sigma(r_B - r_A)$	$r_A + \frac{K}{2} \cdot (\sigma(r_B - r_A) - \sigma(r_A - r_B))$	$r_A - K \cdot \sigma(r_A - r_B)$
$r'_B =$	$r_B - K \cdot \sigma(r_B - r_A)$	$r_B + \frac{K}{2} \cdot (\sigma(r_A - r_B) - \sigma(r_B - r_A))$	$r_B + K \cdot \sigma(r_A - r_B)$

(2)

The motivation for this system comes from thinking of the outcome of a game as a random variable. For some symmetric random variable η , the event $r_A - r_B + \eta > 0$ is recorded as a victory for A , $r_A - r_B + \eta < 0$ as a victory for B , and $r_A - r_B + \eta = 0$ as a draw. Set

$$\sigma(z) = \Pr(\eta < z) + \frac{1}{2} \Pr(\eta = z) \tag{3}$$

to be a kind of ‘symmetrized’ cumulative distribution function of η (which coincides with the usual cumulative distribution function when η has no atoms). Then observe σ satisfies the necessary requirements (1), and that $\mathbb{E}(r'_A) = r_A$ and $\mathbb{E}(r'_B) = r_B$. So one gains points for performing ‘better than expected’ and loses points for performing ‘worse than expected’. Under this probabilistic model, one natural assumption is that η has a finite expectation, though it is not required.

Given some real-world game, one should attempt to pick η so that these estimated probabilities match the empirical win-loss rates observed. The proposal by Elo in the 1960s was to take η to be Gaussian, citing the ubiquity of the normal distribution in nature [1]. However, the community soon decided a logistic random variable was more suitable, leading to the pot function of $\sigma(z) = \frac{1}{1+e^{-cz}}$ for some constant c [2]. The International Chess Federation (FIDE) has long used $c = \log(10)/400 \approx 5.76 \times 10^{-3}$ [4, Chapter B02], though recent analysis suggests $c = \frac{5}{6} \cdot \log(10)/400 \approx 4.80 \times 10^{-3}$ reflects real-world chess data much better [6]. Depending on various factors, FIDE uses K in the range of 10 to 40. The range of ratings exhibited by human players is roughly 0 to 3000 [4, 3]. This paper considers a generic pot function σ , and applies the results to several specific families of pot functions listed in Table 1.

There are additional complications in real-world implementations of Elo. For legibility and practicality, fractional and negative rating points are avoided by scaling and shifting points up and rounding to the nearest integer, and by imposing an artificial floor on possible ratings (by gifting a player points if they would otherwise dip below the floor). The total size of the pot K may also vary depending on various factors, such as how many games each player has played before. For example, K may be large for a new player to facilitate faster convergence of their rating to their true skill level, and may decrease over time to reduce arbitrary fluctuations for experienced players. Sometimes rating updates are batched. That is, one accumulates their pot winnings and losses over several games, and updates their rating once at the end. Often times all the games played at a single tournament are batched together in that manner. All these details and even more complications are outlined thoroughly by different organizations implementing Elo; one may read about them in, for example, the FIDE Handbook [4].

1.1 Setting and results

This paper considers n players starting with equal ratings. Players’ ratings update on a game-by-game basis according to (2) with $K = 1$ kept fixed. In particular, both fractional and negative points are possible and rating updates are not batched. Since the dynamics depend only on the difference between ratings, we can take everyone’s initial rating to be 0 without loss of generality. Asymptotically in k , we seek the highest one of the n players may be rated after a total of k games are played amongst all of them. This question is interesting both for fixed n , and for n allowed to grow with k . We find a phase transition at $n = k^{1/3}$: there is a huge increase in the highest possible Elo rating from $n = 2$ to $n = k^{1/3}$, but (depending on σ) little-to-no increase for any higher n . This paper is organized into four main sections.

Section 2: the highest Elo problem for $n = 2$ is resolved up to constant additive error.

Section 3: a lower bound is provided for each n by finding a family of strategies which achieve a highest rating of $\Theta(\min(n, k^{1/3}))$ for any pot function. Note the bound is saturated for $n = \Omega(k^{1/3})$.

Section 4: an upper bound on the highest possible rating for each n is provided with a mild natural assumption on σ .

Section 5: open questions and a surprising connection to the maximum overhang problem are discussed.

The gap between the upper and lower bounds depends on which σ is used. In particular, the key quantity is the left-tail behaviour of σ . A quickly decaying tail corresponds both to a small asymptotic rate in the case of $n = 2$, and to nearly matching upper and lower bounds for general n . Our main result is stated in terms of the following function, which quantifies the rate of decay of the left-tail of σ :

$$f(x) = \int_0^x \frac{1}{\sigma(-\tau)} d\tau \quad (4)$$

Note that faster the left tail of σ decays to 0, the faster f diverges to infinity. Note that positivity and monotonicity of σ means that f is increasing and convex. This main theorem summarizes the results from each section.

Theorem 1.1. *Let $R(n, k)$ be the highest possible Elo rating achievable with k games n players starting with 0 Elo points. Fix any pot function σ and let f be as in (4). Then*

$$\frac{1}{2}f^{-1}(2k) \leq R(2, k) \leq \frac{1}{2}f^{-1}(2k - 2) + \frac{5}{2}.$$

Now assume that $\sup_z z \sigma(-z) < \infty$. Fix $n = h(k)$. If $h(k) = \Omega(k^{1/3})$, then there exists constants C_1, C_2 such that for sufficiently large k , one has

$$C_1 k^{1/3} \leq R(h(k), k) \leq C_2 k^{1/3} f^{-1}(k)^{1/3}.$$

Furthermore, if $h(k) = o(k^{1/3})$ then

$$C_1 h(k) \leq R(h(k), k) \leq C_2 h(k) f^{-1}(k/h(k)).$$

Proof. $R(2, k)$ handled in Section 2. The lower bounds for $R(h(k), k)$ are in Section 3 and the upper bounds are in Section 4. \square

Remark. The results hold even if the players all start with distinct ratings, provided that all the initial ratings fall into some bounded interval independent of n and k , and that there's some point around which the initial ratings are symmetric. That is, there exists some value r_0 such that for every player rated r there's one rated $2r_0 - r$.

Table 1 lists the value of f^{-1} for several natural families of σ and names the corresponding η where appropriate. For the logistic pot function, f^{-1} is logarithmic and Theorem 1.1 implies

$$R(2, k) = \frac{1}{2c} \log(2k) + O(1) \quad \text{and} \quad R(n, k) = \tilde{\Theta}\left(\min\left(n, k^{1/3}\right)\right)$$

where $\tilde{\Theta}$ suppresses log factors. For any pot function that eventually hits 0 (for example, the 'uniform' pot function $\max(0, \min(1, cz + 1/2))$), let x be such that $\sigma(-x) = 0$ and note f has a vertical asymptote at x . This means f^{-1} is bounded, so Theorem 1.1 implies

$$R(n, k) = \Theta\left(\min\left(n, k^{1/3}\right)\right)$$

where $\tilde{\Theta}$ only suppresses constant factors. At the other extreme, σ need not converge to 0 at all. In this case, $f^{-1}(k) = \Theta(k)$. The upper bounds in Theorem 1.1 no longer hold, but just by noticing the total pot size for each game is bounded, one sees $R(n, k) = \Theta(f^{-1}(k)) = \Theta(k)$ anyway.

η	$\sigma(z)$	$f^{-1}(k)$
Logistic	$\frac{1}{1+e^{-cz}}$	$\frac{1}{c} \log(ck) \cdot (1 + O(1/k))$
–	$\frac{1}{2} \frac{cz}{(1+ cz ^p)^{1/p}} + \frac{1}{2}$	$\left(\frac{1}{2c^p} \cdot \frac{p+1}{p}\right)^{\frac{1}{p+1}} \cdot k^{\frac{1}{p+1}} \cdot (1 + o(1))$
Gaussian	$\frac{1}{2} \operatorname{erf}(cz/\sqrt{2}) + \frac{1}{2}$	$\frac{1}{c} \sqrt{\log k} \cdot (1 + o(1))$
Uniform	$\min\left(1, \max\left(0, cz + \frac{1}{2}\right)\right)$	$\frac{1}{2c} - \frac{e^{-ck}}{2c}$
Cauchy	$\frac{1}{\pi} \arctan(cz) + \frac{1}{2}$	$\sqrt{2/c\pi} \cdot k^{1/2} + O(1)$
–	$\frac{c}{2} \operatorname{sign}(x) + \frac{1}{2}$	$\frac{2}{1-c} \cdot k$

Table 1: The value of f^{-1} for some selected families of pot functions. Most of the families are parameterized by some constant c , which correlates with the slope of each σ at 0. The logistic pot function in the top row is the usual one used by real-world implementations of Elo.

2 Case of $n = 2$

When there are only two players, H and L , as their ratings r_H and r_L grow further apart, the fewer rating points the higher player H can earn from the lower L each time. If $\sigma(-x) = 0$ for some x , the difference in the ratings of H and L cannot exceed $x + 2$. To see this, observe that a game can only increase the value of $r_H - r_L$ if $|r_H - r_L| < x$, and furthermore, r_H and r_L can each change by at most one point in each game. Since $r_H + r_L = 0$ is invariant, this immediately implies

$$r_H \leq \frac{x}{2} + 1$$

independently of how many games are played. On the other hand, if $\sigma(-z) > 0$ for all z , then $r_H - r_L$ will diverge. To see this first note monotonicity of σ means σ is bounded away from 0 on each compact interval. Then, note H beating L will change the value of $r_H - r_L$ by applying the map $g : z \mapsto z + 2\sigma(-z)$. The orbit can only converge if $\sigma(-g^j(z))$ converges to 0, which cannot happen for bounded z by assumption. In the former case, we may still ask how quickly the bound $r_H \leq \frac{x}{2} + 1$ is achieved, and in the latter case we ask how quickly the orbit diverges. Intuitively, the faster $\sigma(-z)$ decays to 0 as $z \rightarrow \infty$, the slower the rate. Indeed, as the next theorem makes explicit, there is a simple expression for the rate in terms of the f defined in (4). Note f is only well defined for $\{x : \sigma(-x) > 0\}$. Nevertheless, it's continuous and strictly increasing, and its range is all of \mathbb{R} . This implies for any σ that it has a well defined inverse f^{-1} defined on all of \mathbb{R} . The following theorem takes advantage of that fact to unify the analysis for every σ .

Theorem 2.1. *Fix any pot function σ . Let $r(t)$ be the highest possible rating after t games with two players. Then*

$$\frac{f^{-1}(2t)}{2} \leq r(t) \leq \frac{f^{-1}(2t-2)}{2} + \frac{5}{2}.$$

Proof. Since the ratings of the two players sum to zero, it suffices to keep track of just the higher rated player. Call the players H and L for ‘higher’ and ‘lower’. Let $r(t)$ be the rating of H after t games. Then $-r(t)$ is the rating of L after t games. Under the assumption that H wins every game, we have the simple recurrence

$$\begin{aligned} r(0) &= 0 & r(t+1) &= r(t) + \sigma(-r(t) - r(t)) \\ & & &= r(t) + \sigma(-2r(t)). \end{aligned}$$

This can be viewed as running Euler’s method on the differential equation

$$y(0) = 0 \quad y' = \sigma(-2y)$$

with a step size of 1. Since y' is positive and monotonically decreasing in y , the Euler approximation upper bounds the exact solution. That is,

$$r(t) \geq y(t).$$

Using separate and integrate, one sees the exact solution to the differential equation is

$$y(t) = \frac{1}{2}f^{-1}(2t).$$

This establishes the lower bound. For the upper bound, we cannot assume that the optimal strategy is for H to win every game. In particular, the function $z \mapsto z + \sigma(-2z)$ need not be monotone, so we cannot exclude the possibility one can achieve a higher rating by first losing a game and ‘slingshotting’ to a higher rating exploiting the fact that $\sigma(-2y)$ is larger for smaller y . Construct a sequence x_j based on the recurrence

$$x_0 = 1 \quad x_{j+1} = x_j + \sigma(-2x_j + 2). \quad (5)$$

This sequence defines a partition of the positive real line, $(0, x_0], (x_0, x_1], (x_1, x_2], (x_2, x_3] \dots$. Again we take $r(t)$ to be the rating of H after t games, but instead of a strict recurrence we only have an inequality $r(t+1) \leq r(t) + \sigma(-2r(t))$, which is tight only when H wins game $t+1$. We now relate $r(t)$ to x_j . If $r(t) \leq 0$, then $r(t+1) \leq r(t) + 1 \leq 1 = x_0$. If $r(t) \in (0, x_0]$, then $r(t+1) \leq r(t) + \sigma(-2r(t)) \leq x_0 + \sigma(-2x_0 + 2) = x_1$. If $r(t) \in (x_{j-1}, x_j]$ for some j , then

$$\begin{aligned} r(t+1) &\leq r(t) + \sigma(-2r(t)) \\ &\leq x_j + \sigma(-2x_{j-1}) \\ &\leq x_j + \sigma(-2x_j + 2) \\ &= x_{j+1}. \end{aligned}$$

In each case, we see that

$$r(t) \leq x_j \implies r(t+1) \leq x_{j+1}.$$

Since $r(0) = 0 < 1 = x_0$, we have $r(t) \leq x_t$ regardless of the sequence of wins and losses. Now similar to before, x_t can be viewed as the result of running Euler’s method with a step size of 1 on the differential equation

$$y' = \sigma(-2y + 2). \quad (6)$$

Using separate and integrate, see that the solution to the differential equation is

$$y(t) = \frac{1}{2}f^{-1}(2t) + 1.$$

Again the Euler approximation is an upper bound to the exact solution. That is,

$$x_t \geq y(t). \quad (7)$$

Suprisingly, this lower bound leads us to an upper bound. We first plug the recurrence (5) back into itself to express x_t as a sum, then use (7) with monotonicity of σ .

$$x_t = x_0 + \sum_{j=0}^{t-1} \sigma(-2x_j + 2) \leq x_0 + \sum_{j=0}^{t-1} \sigma(-2y(j) + 2). \quad (8)$$

The summand is the same as the right hand side of the differential equation (6). Then the sum is Reimann sum of step size 1. But since the summand is monotone, we can bound the Reimann sum by the integral.

$$\sum_{j=0}^{t-1} \sigma(-2y(j) + 2) = \sum_{j=0}^{t-1} y'(j) \leq y'(0) + \int_0^{t-1} y'(s) ds = y'(0) + y(t-1). \quad (9)$$

Note $y'(0) = 1/2$ and $x_0 = 1$, so combining (8) and (9) gives the final upper bound of

$$x_t \leq 1 + \frac{1}{2} + y(t-1) = 1 + \frac{1}{2} + 1 + \frac{1}{2}f^{-1}(2t-2).$$

□

3 Lower bound for general n

This section describes two strategies for any number of players and games. The first strategy does not depend at all on the pot function used, but requires that everyone's initial ratings be exactly equal. The second strategy has a small dependence on the pot function used, but works for any symmetric list of initial ratings. Both strategies produce higher ratings for larger n , up to $n = \Theta(k^{1/3})$, at which point the asymptotic highest rating in k is $\Theta(k^{1/3})$. The first strategy is very simple: **pick any pair of players with equal rating and have one beat the other. Repeat until all players have a distinct rating or k games have been played.** This strategy is guaranteed to produce a player of either very high or very low rating. If it produces a player of very low rating, simply re-do the strategy picking the same sequence of pairs of players but have the opposite player win. Since game outcomes are symmetric, this will produce a player of high rating instead.

Theorem 3.1. *The aforementioned strategy achieves a highest rating of*

$$\min\left(\frac{k^{1/3}}{2}, \frac{n}{4}\right)$$

when all players start with 0 rating points. In particular, the rating achieved is $\Omega(k^{1/3})$ for $n = \Omega(k^{1/3})$.

Proof. Let $\mathbf{r}(j)$ be the multiset of player ratings after j games are played and $|\mathbf{r}(j)|$ its element absolute value. We claim that every entry of $\mathbf{r}(j)$ is a half-integer. To see this, note $\sigma(0) = 1/2$ for every pot function so player ratings change in increments of $1/2$ and all start at 0. We next claim

$$\sum_{r \in \mathbf{r}(j)} r^2 = j/2. \tag{10}$$

We argue inductively. The claim is clearly true initially when all ratings are 0. When two players rated a points play, the increase in the functional is

$$\Delta = (a + 1/2)^2 + (a - 1/2)^2 - a^2 - a^2 = 1/2.$$

Suppose the strategy terminates when k games have been played. Then

$$\frac{k}{2} = \sum_{r \in \mathbf{r}(k)} r^2 \leq n \max |\mathbf{r}(k)|^2 \implies \max |\mathbf{r}(k)| \geq \sqrt{k/2n}.$$

Now suppose the strategy terminates when all players have distinct ratings. Then since all ratings are half integers, that means $\max |\mathbf{r}(j)| \geq n/4$. Without a guarantee on which way the strategy terminates, we get the worst of both bounds. So the maximum rating is at least $\min\left(\sqrt{k/2n}, n/4\right)$. If $n \leq 2k^{1/3}$, the minimum is $n/4$. If $n > 2k^{1/3}$, we can rerun the strategy ignoring all but the first $2k^{1/3}$ players, giving a bound of exactly $k^{1/3}/2$. \square

Remark. Theorem 3.1 is unfortunately very brittle. If the initial ratings are perturbed slightly, we can no longer assume all ratings are half-integers, nor can we often expect players to have equal rating. For

$$\sigma(x) = \frac{1}{2} \text{sign}(x) + \frac{1}{2} = \begin{cases} 1 & x > 0 \\ 1/2 & x = 0 \\ 0 & x < 0 \end{cases},$$

this is indeed the best we can do. If all players have slightly different ratings, no Elo can be transferred between them. However, there's a robust version of the strategy that works for every other σ : first fix some $\delta > 0$ for which $\sigma(-\delta) > 0$. **Pick any pair of players whose ratings are within δ of each other**

and have the higher rated player beat the lower rated player. Repeat until no two players are within δ rating points or until k games have been played. Note for $\delta = 0$, one recovers the original strategy. As before, we may have to flip everyone's rating to ensure we end with a very high rating, as opposed to a very low rating. We need not assume that the initial ratings are all 0 anymore, but we do need to allow ourselves this possible reflection.

Theorem 3.2. Fix any $\vec{r} = (r_1, \dots, r_n)$. The aforementioned strategy achieves a highest rating of

$$\min\left(\delta^{1/3}\sigma(-\delta)^{1/3} \cdot k^{1/3}, \frac{\delta}{2} \cdot n\right)$$

for initial ratings $\mathbf{r}(0) = \vec{r}$ or $\mathbf{r}(0) = -\vec{r}$.

Proof. The proof is very similar to the proof of Theorem 3.1. Again we consider the functional $\sum_{r \in \mathbf{r}(j)} r^2$. In a game between players rated $a < b$, the change in the functional is

$$\begin{aligned} \Delta &= (b + \sigma(a - b))^2 + (a - \sigma(a - b))^2 - b^2 - a^2 \\ &= 2b\sigma(a - b) - 2a\sigma(a - b) + 2\sigma(a - b)^2 \\ &= 2(b - a)\sigma(a - b) + 2\sigma(a - b)^2. \\ &\geq 2\sigma(-\delta)^2. \end{aligned}$$

Suppose the strategy terminates when k games have been played. Then

$$2\sigma(-\delta) \cdot k \leq \sum_{r \in \mathbf{r}(k)} r^2 \leq n \max |\mathbf{r}(k)|^2 \implies \max |\mathbf{r}(k)| \geq \sqrt{2\sigma(-\delta)k/n}.$$

Now suppose the strategy terminates when no two players are within δ rating points. Since there are a total of n players, that means $\max |\mathbf{r}(j)| \geq \delta n/2$. As before, without a guarantee on which way the strategy terminates, we get the worst of both bounds. When

$$n \leq \frac{2\sigma(-\delta)^{1/3}}{\delta^{2/3}} \cdot k^{1/3},$$

the smaller bound is $\delta n/2$. When n is larger, one simply ignores the excess players, as before. \square

Remark. The theorem is best applied with the vector of initial ratings is symmetric, i.e. \vec{r} is a permutation of $-\vec{r}$. In fact, by shifting all the points up or down by a constant amount, they can be symmetric around any (constant valued) point, i.e. \vec{r} is a permutation of $(r_0, \dots, r_0) - \vec{r}$.

Remark. This bound is stronger the heavier the tail of σ . Consider, for instance $\sigma(z) = \frac{1}{1+e^{-cz}}$ as $c \rightarrow \infty$, or any other family of σ approaching the pathological $\frac{1}{2}\text{sign}(x) + \frac{1}{2}$. The bound of Theorem 3.2 becomes weaker and weaker. This contrasts with the bound of Theorem 3.1, which is completely independent of the pot function but requires initial ratings be exactly equal.

4 Upper bound for general n

In this section, we wish to show the algorithm presented in Section 3 is nearly optimal. Our strategy is to show that achieving a rating of r requires many games to be played. We start with a relaxation of the setup: instead of considering a discrete sequence of games resulting in a discrete sequence of player ratings, we consider a continuous path in the space of possible player ratings. Call a path $\mathbf{r} : [0, c] \rightarrow \mathbb{R}^n$ **valid** if there exists a finite sequence $0 = t_0 < \dots < t_k = c$ such that for each $j \in [k]$ there exists $w, \ell \in [n]$ with

$$\mathbf{r}'(t) = \mathbf{e}_w - \mathbf{e}_\ell \quad \forall t \in (t_{j-1}, t_j)$$

where $\mathbf{e}_w, \mathbf{e}_\ell$ are elementary basis vectors. For convenience, we write $\mathbf{r}'(t_{j-1}) = \mathbf{e}_w - \mathbf{e}_\ell$ to make \mathbf{r}' right-continuous. In other words, \mathbf{r} has constant speed and consists of line segments (called **edges**) along which only two coordinates change. The points $\mathbf{r}(t_j)$ are referred to as **vertices**. We will use the notation

$$A \rightarrow B \rightarrow C$$

to denote the path with vertices A, B, C . A sequence of games corresponds naturally to a valid path: let $\mathbf{r}(t_j)$ be the list of the players' ratings after game j . The indices ' w ' and ' ℓ ' conveniently correspond to the 'winner' and 'loser' for each game. In the event of a draw, the initially lower rated player is considered the winner as their rating increases. Note

$$t_{j+1} - t_j = \text{amount of Elo transferred for game } j + 1. \quad (11)$$

Define a weight function

$$W(\mathbf{x}, \mathbf{y}) = \begin{cases} \frac{1}{\sigma(2 - \langle \mathbf{x}, \mathbf{y} \rangle)} & \langle \mathbf{x}, \mathbf{y} \rangle > 0 \\ 0 & \langle \mathbf{x}, \mathbf{y} \rangle \leq 0 \end{cases} \quad (12)$$

and define the length of the path to be

$$\text{len}(\mathbf{r}) = \int_0^c W(\mathbf{r}'(t), \mathbf{r}(t)) dt.$$

Whereas the Euclidean length of a path is simply $c\sqrt{2}$, the value of $\text{len}(\mathbf{r})$ depends highly on the motion of the path. When a valid path comes from a sequence of games, Lemma 4.1 shows $\text{len}(\mathbf{r})$ cannot be larger than k . The intuition is that $\langle \mathbf{r}'(t), \mathbf{r}(t) \rangle = \mathbf{r}_w(t) - \mathbf{r}_\ell(t)$ records the difference in two players' ratings, so the Euclidean length of the $(j + 1)$ th edge is $\sqrt{2}\sigma(-\langle \mathbf{r}'(t_j), \mathbf{r}(t_j) \rangle)$, which approximately cancels with $w(\mathbf{r}'(t), \mathbf{r}(t))$ resulting in constant contribution.

Lemma 4.1. *Let \mathbf{r} be a valid path arising from a sequence of k games. Then*

$$\text{len}(\mathbf{r}) \leq k.$$

Proof. It suffices to show for each edge $\mathbf{r}(t_j) \rightarrow \mathbf{r}(t_{j+1})$ that

$$\text{len}(\mathbf{r}(t_j) \rightarrow \mathbf{r}(t_{j+1})) \leq 1$$

since summing over all edges would produce the final result. Fix any j . Let w, ℓ be the indices such that

$$\mathbf{r}'(t) = \mathbf{e}_w - \mathbf{e}_\ell \quad \forall t \in [t_j, t_{j+1}).$$

Let $z = \langle \mathbf{r}'(t_j), \mathbf{r}(t_j) \rangle$ so that by (11) we have

$$\sigma(-z) = t_{j+1} - t_j.$$

Also note $\langle \mathbf{r}'(t), \mathbf{r}(t) \rangle = z + 2(t - t_j) \leq z + 2\sigma(-z) \leq z + 2$. Then

$$\begin{aligned} \text{len}(\mathbf{r}(t_j) \rightarrow \mathbf{r}(t_{j+1})) &= \int_{t_j}^{t_{j+1}} W(\mathbf{r}'(t), \mathbf{r}(t)) dt \\ &\leq \int_{t_j}^{t_{j+1}} \frac{1}{\sigma(2 - \langle \mathbf{r}'(t), \mathbf{r}(t) \rangle)} dt \\ &\leq \int_{t_j}^{t_{j+1}} \frac{1}{\sigma(-z)} dt \\ &= 1. \end{aligned}$$

□

Lemma 4.1 means that instead of bounding k , we can bound $\text{len}(\mathbf{r})$. Directly doing so for any valid path is difficult, so we introduce two additional notions that allow us to make some additional restrictions on \mathbf{r} . Call a valid path **ordered** if

$$\mathbf{r}_1(t) \geq \dots \geq \mathbf{r}_n(t) \quad \forall t \in [0, c].$$

Each valid path corresponds to an ordered valid path of the same length by doing the following: each time \mathbf{r} intersects a hyperplane of the form $x_w = x_\ell$, reflect the remaining part of the path across it. This makes it possible to refer unambiguously to the p th highest rated player. It also means $\langle \mathbf{r}', \mathbf{r} \rangle$ does not change sign on each edge. The last notion we introduce is that of an **upset**. This is when a lower rated player beats or draws with a higher rated player. Stated in terms of valid paths, an upset is an edge for which

$$\langle \mathbf{r}', \mathbf{r} \rangle = \mathbf{r}_w - \mathbf{r}_\ell < 0$$

at the starting vertex. We essentially show an optimal strategy, when converted to a ordered valid path, does not make use of upsets. The precise statement is given in Lemma 4.2. The intuition here is that upsets bring the players' ratings closer together, whereas in order to make one rating large you need the ratings to be spread out. The proof strategy is to take any valid path with upsets and convert it to one achieving a higher maximum rating without upsets. This lemma is the longest and most technical as it involves some casework in describing this conversion procedure.

Lemma 4.2. *For each valid ordered path $\mathbf{r} : [0, c] \rightarrow \mathbb{R}^n$ there exists a valid ordered path $\tilde{\mathbf{r}} : [0, \tilde{c}] \rightarrow \mathbb{R}^n$ with*

$$\text{len}(\tilde{\mathbf{r}}) \leq \text{len}(\mathbf{r}) \quad \text{and} \quad \max \tilde{\mathbf{r}}(\tilde{c}) \geq \max \mathbf{r}(c)$$

such that \mathbf{r}' has no upsets.

Proof. Given \mathbf{r} , we start by constructing $\tilde{\mathbf{r}}$ with a slightly different property than the requirement of the theorem. We require $\mathbf{r}(c) = \tilde{\mathbf{r}}(\tilde{c})$ and that all the upsets in $\tilde{\mathbf{r}}$ occur at the end. That is, once $\langle \tilde{\mathbf{r}}', \tilde{\mathbf{r}} \rangle$ becomes negative, it stays negative. Once we have done that, observe that upsets only decrease the value of $\max \tilde{\mathbf{r}}$ so we obtain a larger maximum value by truncating the path just before the upsets.

Fix vector $\mathbf{u} = (1, 2, \dots, n)$. Then each valid path \mathbf{r} has an associated sequence $S(\mathbf{r})$ of integers $S(\mathbf{r})_j = \langle \mathbf{u}, \mathbf{r}'(t_j) \rangle$ with one integer for each edge. Since that path is ordered, each integer is positive if and only if the corresponding edge is an upset. Equip the set of possible values of $S(\mathbf{r})$ with the lexicographic ordering. Our method for constructing $\tilde{\mathbf{r}}$ is iterative, where each iteration strictly decreases $S(\mathbf{r})$ and does not increase $\text{len}(\mathbf{r})$. The strict decreasing of $S(\mathbf{r})$ guarantees that this procedure terminates after at most $n^\#$ edges in \mathbf{r} steps.

We now describe one iteration. Given a path \mathbf{r} , locate three consecutive vertices $\mathbf{r}(t_j), \mathbf{r}(t_{j+1}), \mathbf{r}(t_{j+2})$ with

$$S(\mathbf{r})_j > 0 \quad \text{and} \quad S(\mathbf{r})_{j+1} < 0.$$

That is, the two connecting edges are a upset followed by a non-upset. We modify the path to obtain $\tilde{\mathbf{r}}$ by either deleting the middle vertex $\mathbf{r}(t_{j+1})$ from the path and adding an edge directly from $\mathbf{r}(t_j)$ to $\mathbf{r}(t_{j+2})$, or replacing the middle vertex with new vertex \mathbf{s} . If $\mathbf{r}(t_{j+1})$ is deleted, $S(\mathbf{r})$ is shortened and therefore decreased. If $\mathbf{r}(t_{j+1})$ is replaced by \mathbf{s} , note that only the j th and $(j+1)$ th entries in $S(\mathbf{r})$ are affected. So requiring that

$$S(\tilde{\mathbf{r}})_j < S(\mathbf{r})_j$$

suffices for $S(\mathbf{r})$ to decrease overall. Note the right hand side is positive, so it suffices for $S(\tilde{\mathbf{r}})_j$ to be negative. We additionally require

$$\text{len}(\tilde{\mathbf{r}}) \leq \text{len}(\mathbf{r}).$$

The modification differs based on how many coordinates change on those two edges. Edge edge modifies two coordinates, so this number can be 2, 3, or 4.

Case of 4 coordinates: this corresponds to the players of the two games being disjoint pairs. Intuitively, the order of the games is irrelevant to the outcome. Set $\mathbf{s} = \mathbf{r}(t_j) - \mathbf{r}(t_{j+1}) + \mathbf{r}(t_{j+2})$. Then $S(\tilde{\mathbf{r}})_j =$

$S(\mathbf{r})_{j+1} < 0$ and $\text{len}(\mathbf{r}) = \text{len}(\tilde{\mathbf{r}})$. **Case of 2 coordinates:** this corresponds to the same two players playing in consecutive games. Let those players be w, ℓ with $w < \ell$. Let Π be the projection onto the w, ℓ coordinates. Then for some $a, b > 0$,

$$\Pi\mathbf{r}(t_j) = \begin{bmatrix} x \\ y \end{bmatrix} \quad \Pi\mathbf{r}(t_{j+1}) = \begin{bmatrix} x - a \\ y + a \end{bmatrix} \quad \Pi\mathbf{r}(t_{j+2}) = \begin{bmatrix} x - a + b \\ y + a - b \end{bmatrix}.$$

Our modification simply deletes $\mathbf{r}(t_{j+1})$, so S certainly decreases. The original length is

$$\begin{aligned} \text{len}\left(\begin{bmatrix} x \\ y \end{bmatrix} \rightarrow \begin{bmatrix} x - a \\ y + a \end{bmatrix} \rightarrow \begin{bmatrix} x - a + b \\ y + a - b \end{bmatrix}\right) &= \int_0^b \frac{1}{\sigma(2 + (y + a - t) - (x - a + t))} dt \\ &= \int_0^b \frac{1}{\sigma(2 + y - x + 2a - 2t)} dt. \\ &= \int_{-a}^{b-a} \frac{1}{\sigma(2 + y - x - 2t)} dt. \end{aligned}$$

If $a \geq b$, the new length is 0. If $a < b$, then

$$\begin{aligned} \text{len}\left(\begin{bmatrix} x \\ y \end{bmatrix} \rightarrow \begin{bmatrix} x - a + b \\ y + a - b \end{bmatrix}\right) &= \int_0^{b-a} \frac{1}{\sigma(2 + (y - t) - (x + t))} dt \\ &= \int_0^{b-a} \frac{1}{\sigma(2 + y - x - 2t)} dt \end{aligned}$$

which is strictly less than the original length since the integrand is the same with a smaller range. **Case of 3 coordinates:** this corresponds to one player playing two different opponents. Let Π be the projection onto those three coordinates in ranked order, i.e.

$$\Pi\mathbf{r}(t_j) = \begin{bmatrix} x \\ y \\ z \end{bmatrix}$$

for $x \geq y \geq z$. There are six possible sign patterns of $\Pi(\mathbf{r}(t_{j+2}) - \mathbf{r}(t_j))$, each with a different selection of \mathbf{s} .

Subcase (+, -, +): For some $a, c > 0$ we have

$$\Pi\mathbf{r}(t_{j+2}) = \begin{bmatrix} x + a \\ y - a - c \\ z + c \end{bmatrix}.$$

There are three possible values of $\Pi\mathbf{r}(t_{j+1})$ corresponding to the three possible upsets that can occur. No matter which one we observe, the selection of their replacement \mathbf{s} is the same.

$$\Pi\mathbf{r}(t_{j+1}) = \begin{bmatrix} x \\ y - c \\ z + c \end{bmatrix}, \begin{bmatrix} x \\ y - a - c \\ z + a + c \end{bmatrix}, \begin{bmatrix} x - c \\ y \\ z + c \end{bmatrix}, \quad \mathbf{s} = \begin{bmatrix} x + a \\ y - a \\ z \end{bmatrix}$$

The possible original lengths are as follows.

$$\begin{aligned}
L_1 &= \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x \\ y-c \\ z+c \end{bmatrix} \rightarrow \begin{bmatrix} x+a \\ y-a-c \\ z+c \end{bmatrix} \right) = \int_0^a \frac{1}{\sigma(2 + (y-c-t) - (x+t))} dt \\
&= \int_0^a \frac{1}{\sigma(2 + y - x - c - 2t)} dt, \\
L_2 &= \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x \\ y-a-c \\ z+a+c \end{bmatrix} \rightarrow \begin{bmatrix} x+a \\ y-a-c \\ z+c \end{bmatrix} \right) = \int_0^a \frac{1}{\sigma(2 + (z+a+c-t) - (x+t))} dt \\
&= \int_0^a \frac{1}{\sigma(2 + z - x + a + c - 2t)} dt, \\
L_3 &= \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x-c \\ y \\ z+c \end{bmatrix} \rightarrow \begin{bmatrix} x+a \\ y-a-c \\ z+c \end{bmatrix} \right) = \int_0^{a+c} \frac{1}{\sigma(2 + (y-t) - (x-c+t))} dt \\
&= \int_0^{a+c} \frac{1}{\sigma(2 + y - x + c - 2t)} dt \\
&= \int_{-c}^a \frac{1}{\sigma(2 + y - x - c - 2t)} dt
\end{aligned}$$

The new length is

$$L_4 = \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x+a \\ y-a \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x+a \\ y-a-c \\ z+c \end{bmatrix} \right) = \int_0^a \frac{1}{\sigma(2 + (y-t) - (x+t))} dt = \int_0^a \frac{1}{\sigma(2 + y - x - 2t)} dt.$$

By monotonicity of σ , we automatically have $L_3 \geq L_1 \geq L_4$. We assume \mathbf{r} is an ordered path, so in the L_2 case we have $z+a+c \leq y-a-c$ giving $L_2 \geq L_4$ showing our modification did not increase the length. Also see that $\mathbf{r}(t_j) \rightarrow \mathbf{s}$ is not an upset so $S(\tilde{\mathbf{r}})_j < 0$.

Subcase $(-, +, -)$: For some $a, c > 0$ we have

$$\Pi\mathbf{r}(t_{j+2}) = \begin{bmatrix} x-a \\ y+a+c \\ z-c \end{bmatrix}.$$

Again there are three possible values of $\Pi\mathbf{r}(t_{j+1})$ corresponding to the three possible upsets that can occur, and our selection of their replacement \mathbf{s} is the same regardless.

$$\Pi\mathbf{r}(t_{j+1}) = \begin{bmatrix} x-a \\ y+a \\ z \end{bmatrix}, \begin{bmatrix} x-a-c \\ y+a+c \\ z \end{bmatrix}, \begin{bmatrix} x-a \\ y \\ z+a \end{bmatrix}, \quad \mathbf{s} = \begin{bmatrix} x \\ y+c \\ z-c \end{bmatrix}.$$

The original lengths are as follows.

$$\begin{aligned}
L_1 &= \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x-a \\ y+a \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x-a \\ y+a+c \\ z-c \end{bmatrix} \right) = \int_0^c \frac{1}{\sigma(2+(z-t)-(y+a+t))} dt \\
&= \int_0^c \frac{1}{\sigma(2+z-y-a-2t)} dt, \\
L_2 &= \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x-a-c \\ y+a+c \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x-a \\ y+a+c \\ z-c \end{bmatrix} \right) = \int_0^c \frac{1}{\sigma(2+(z-t)-(x-a-c+t))} dt \\
&= \int_0^c \frac{1}{\sigma(2+z-x+a+c-2t)} dt, \\
L_3 &= \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x-a \\ y \\ z+a \end{bmatrix} \rightarrow \begin{bmatrix} x-a \\ y+a+c \\ z-c \end{bmatrix} \right) = \int_0^{a+c} \frac{1}{\sigma(2+(z-t)-(y+t))} dt \\
&= \int_0^{a+c} \frac{1}{\sigma(2+z-y-2t)} dt, \\
&= \int_{-a}^c \frac{1}{\sigma(2+z-y-2a-2t)} dt.
\end{aligned}$$

The new length is

$$L_4 = \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x \\ y+c \\ z-c \end{bmatrix} \rightarrow \begin{bmatrix} x-a \\ y+a+c \\ z-c \end{bmatrix} \right) = \int_0^c \frac{1}{\sigma(2+(z-t)-(y+t))} dt = \int_0^c \frac{1}{\sigma(2+z-y-2t)} dt.$$

Again by monotonicity of σ , we automatically have $L_3 \geq L_1 \geq L_4$. We assume \mathbf{r} is an ordered path, so in the L_2 case we have $y+a+c \leq x-a-c$ giving $L_2 \geq L_4$ as required. Also see that $\mathbf{r}(t_j) \rightarrow \mathbf{s}$ is not an upset so $S(\tilde{\mathbf{r}})_j < 0$.

Subcase $(+, +, -)$: For some $a, b > 0$ we have

$$\Pi\mathbf{r}(t_{j+2}) = \begin{bmatrix} x+a \\ y+b \\ z-a-b \end{bmatrix}.$$

There is only one possible value of $\Pi\mathbf{r}(t_2)$ corresponding to an upset.

$$\Pi\mathbf{r}(t_{j+1}) = \begin{bmatrix} x-b \\ y+b \\ z \end{bmatrix}, \quad \mathbf{s} = \begin{bmatrix} x \\ y+b \\ z-b \end{bmatrix}.$$

The original length is

$$\begin{aligned}
L_1 &= \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x-b \\ y+b \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x+a \\ y+b \\ z-a-b \end{bmatrix} \right) = \int_0^{a+b} \frac{1}{\sigma(2+(z-t)-(x-b+t))} dt \\
&= \int_0^{a+b} \frac{1}{\sigma(2+z-x+b-2t)} dt \\
&= \int_0^b \frac{1}{\sigma(2+z-x+b-2t)} dt + \int_0^a \frac{1}{\sigma(2+z-x-b-2t)} dt
\end{aligned}$$

and the new length is

$$\begin{aligned} L_2 &= \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x \\ y+b \\ z-b \end{bmatrix} \rightarrow \begin{bmatrix} x+a \\ y+b \\ z-a-b \end{bmatrix} \right) = \int_0^b \frac{1}{\sigma(2+(z-t)-(y+t))} dt + \int_0^a \frac{1}{\sigma(2+(z-b-t)-(x+t))} dt \\ &= \int_0^b \frac{1}{\sigma(2+z-y-2t)} dt + \int_0^a \frac{1}{\sigma(2+z-x-b-2t)} dt. \end{aligned}$$

Again since we assume the path is ordered, we have $x \geq y+b$ so $L_2 \leq L_1$ by monotonicity of σ . Also see that $\mathbf{r}(t_j) \rightarrow \mathbf{s}$ is not an upset so $S(\tilde{\mathbf{r}})_j < 0$.

Subcase (+, -, -): For some $b, c > 0$, we have

$$\Pi\mathbf{r}(t_{j+2}) = \begin{bmatrix} x+b+c \\ y-b \\ z-c \end{bmatrix}.$$

There is only one possible value of $\Pi\mathbf{r}(t_2)$ corresponding to an upset.

$$\Pi\mathbf{r}(t_{j+1}) = \begin{bmatrix} x \\ y-b \\ z+b \end{bmatrix}, \quad \mathbf{s} = \begin{bmatrix} x+b \\ y-b \\ z \end{bmatrix}.$$

The original length is

$$\begin{aligned} L_1 &= \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x \\ y-b \\ z+b \end{bmatrix} \rightarrow \begin{bmatrix} x+b+c \\ y-b \\ z-c \end{bmatrix} \right) = \int_0^{b+c} \frac{1}{\sigma(2+(z+b-t)-(x+t))} dt \\ &= \int_0^{b+c} \frac{1}{\sigma(2+z-x+b-2t)} dt \\ &= \int_0^b \frac{1}{\sigma(2+z-x+b-2t)} dt + \int_0^c \frac{1}{\sigma(2+z-x-b-2t)} dt \end{aligned}$$

and the new length is

$$\begin{aligned} L_2 &= \text{len} \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x+b \\ y-b \\ z \end{bmatrix} \rightarrow \begin{bmatrix} x+b+c \\ y-b \\ z-c \end{bmatrix} \right) = \int_0^b \frac{1}{\sigma(2+(y-t)-(x+t))} dt + \int_0^c \frac{1}{\sigma(2+(z-t)-(x+b+t))} dt \\ &= \int_0^b \frac{1}{\sigma(2+y-x-2t)} dt + \int_0^c \frac{1}{\sigma(2+z-x-b-2t)} dt. \end{aligned}$$

Since we assume the path is ordered, we have $y-b \geq z$ so $L_2 \leq L_1$ once again. Also see that $\mathbf{r}(t_j) \rightarrow \mathbf{s}$ is not an upset so $S(\tilde{\mathbf{r}})_j < 0$.

Subcase (-, +, +), (-, -, +): For some a, c we have

$$\Pi\mathbf{r}(t_{j+2}) = \begin{bmatrix} x-a \\ y+a-c \\ z+c \end{bmatrix}.$$

Depending on the sign of $a-c$, there are two possible values of $\Pi\mathbf{r}(t_{j+1})$. In either case, the selection of \mathbf{s} is the same.

$$\Pi\mathbf{r}(t_{j+1}) = \begin{bmatrix} x-a \\ y \\ z+a \end{bmatrix}, \begin{bmatrix} x-c \\ y \\ z+c \end{bmatrix} \quad \mathbf{s} = \begin{bmatrix} x-a \\ y+a \\ z \end{bmatrix}.$$

In this case, both new edges are upsets so the new length is 0, but we cannot conclude $S(\tilde{\mathbf{r}})_j < 0$ as before. Instead, let $w_1 < w_2 < w_3$ be the indices of the three relevant coordinates and note

$$S(\mathbf{r})_j = w_1 - w_3 > S(\tilde{\mathbf{r}})_j = w_1 - w_2.$$

Finally note the sign patterns of $(+, +, +)$ and $(-, -, -)$ are not possible since the sum of all ratings is invariant. Since in all possible cases we have a decrease of S without an increase of length, this process terminates in a path that isn't longer and doesn't have any upsets followed by a non-upset. We end by truncating off any upsets at the end of the path. \square

Lemma 4.2 implies we can restrict our attention to valid ordered upset-free paths. The length of these paths can be bounded in terms of the following potential function.

Definition 4.1. Define $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}$ by

$$\Phi(\mathbf{s}) = \|\mathbf{s}\|^2 + \sum_{p=1}^{n-1} f(-2 + \mathbf{s}_p - \mathbf{s}_{p+1})$$

where f is as in (4).

We seek to show Φ grows slowly as games are played, so that Φ of the end point of the path is upper bounded by (a multiple of) the length of the path. Note that the first term $\|\mathbf{s}\|^2$ in this potential function is exactly the $\sum r^2$ expression appearing in the proofs of Theorems 3.1 and 3.2. The intuitive idea behind the second term is that when player p beats player $p+1$, the ratings \mathbf{s}_p and \mathbf{s}_{p+1} move $2\sigma(\mathbf{s}_{p+1} - \mathbf{s}_p)$ apart, which by the definition of f means the corresponding f term increases by just a constant amount. Lemma 4.3 makes this intuition precise.

Lemma 4.3. For a valid ordered upset-free path $\mathbf{r} : [0, c] \rightarrow \mathbb{R}^n$ one has

$$\Phi(\mathbf{r}(c)) - \Phi(\mathbf{r}(0)) \leq \left(2 + 2 \sup_z \sigma(2 - z)\right) \cdot \text{len}(\mathbf{r})$$

Proof. By the fundamental theorem of calculus, it suffices to show that

$$\frac{d}{dt} \Phi(\mathbf{r}(t)) \leq \left(2 + 2 \sup_z \sigma(2 - z)\right) W(\mathbf{r}'(t), \mathbf{r}(t)).$$

Since \mathbf{r} is upset-free we have, $\langle \mathbf{r}'(t), \mathbf{r} \rangle \geq 0$ so by the definition of W (12),

$$W(\mathbf{r}'(t), \mathbf{r}(t)) = \frac{1}{\sigma(2 - \langle \mathbf{r}'(t), \mathbf{r}(t) \rangle)}.$$

The chain rule gives

$$\frac{d}{dt} \Phi(\mathbf{r}(t)) = \langle \mathbf{r}'(t), \nabla \Phi(\mathbf{r}(t)) \rangle. \quad (13)$$

Since \mathbf{r} is ordered in addition to being upset-free, $\mathbf{r}'(t)$ is always of the form $\mathbf{r}'(t) = \mathbf{e}_j - \mathbf{e}_{j+a}$ for positive a . That is, the winning player has a lower index than the losing player. The j th entry of the gradient can be computed directly as

$$\begin{aligned} \frac{\partial}{\partial \mathbf{s}_j} \Phi(\mathbf{s}) &= 2\mathbf{s}_j + f'(-2 + \mathbf{s}_j - \mathbf{s}_{j+1}) \cdot \mathbf{1}_{j < n} - f'(-2 + \mathbf{s}_{j-1} - \mathbf{s}_j) \cdot \mathbf{1}_{j > 1} \\ &= 2\mathbf{r}_j + \frac{\mathbf{1}_{j < n}}{\sigma(2 + \mathbf{s}_{j+1} - \mathbf{s}_j)} - \frac{\mathbf{1}_{j > 1}}{\sigma(2 + \mathbf{s}_j - \mathbf{s}_{j-1})}, \end{aligned}$$

and in particular

$$\begin{aligned}
\langle \mathbf{e}_j - \mathbf{e}_{j+a}, \nabla \Phi(\mathbf{s}) \rangle &= \left(\frac{\partial}{\partial \mathbf{s}_j} \Phi(\mathbf{s}) - \frac{\partial}{\partial \mathbf{s}_{j+a}} \Phi(\mathbf{s}) \right) \leq 2(\mathbf{s}_j - \mathbf{s}_{j+a}) + \frac{\mathbf{1}_{j < n}}{\sigma(2 + \mathbf{s}_{j+1} - \mathbf{s}_j)} + \frac{\mathbf{1}_{j+a > 1}}{\sigma(2 + \mathbf{s}_{j+a} - \mathbf{s}_{j+a-1})} \\
&\leq 2(\mathbf{r}_j - \mathbf{r}_{j+a}) + \frac{\mathbf{1}_{j < n}}{\sigma(2 + \mathbf{s}_{j+a} - \mathbf{s}_j)} + \frac{\mathbf{1}_{j+a > 1}}{\sigma(2 + \mathbf{s}_{j+a} - \mathbf{s}_j)} \\
&\leq \frac{2 \sup_z z \sigma(2 - z)}{\sigma(2 + \mathbf{s}_{j+a} - \mathbf{s}_j)} + \frac{2}{\sigma(2 + \mathbf{s}_{j+a} - \mathbf{s}_j)} \\
&= \left(2 + 2 \sup_z z \sigma(2 - z) \right) \cdot \frac{1}{\sigma(2 - \langle \mathbf{e}_j - \mathbf{e}_{j+a}, \mathbf{s} \rangle)} \\
&= \left(2 + 2 \sup_z z \sigma(2 - z) \right) \cdot w(\mathbf{e}_j - \mathbf{e}_{j+a}, \mathbf{s}).
\end{aligned}$$

For $\mathbf{s} = \mathbf{r}(t)$ and $\mathbf{e}_j - \mathbf{e}_{j+a} = \mathbf{r}'(t)$, this is exactly what we needed to show. \square

The last piece of the puzzle is showing $\Phi(\mathbf{s})$ has to be large whenever one entry of \mathbf{s} is large. Rapid growth of f means if any two consecutive players have a large rating difference, Φ will be large. On the other hand, if consecutive players are close in rating, many entries of \mathbf{s} are close to its largest entry forcing $\|\mathbf{s}\|^2$ to be large.

Lemma 4.4. *Let \mathbf{s} be any vector such that $\max \mathbf{s} = \mathbf{s}_1$ and at least one entry is negative. Then*

$$\Phi(\mathbf{s}) \geq \mathbf{s}_1^2 \quad \text{and} \quad \Phi(\mathbf{s}) \geq \frac{\mathbf{s}_1^3/8}{f^{-1}(\mathbf{s}_1^2/4) + 2}.$$

Furthermore,

$$\Phi(\mathbf{s}) \geq n f\left(-2 + \frac{\mathbf{s}_1}{2n}\right).$$

Proof. The first result is immediate from $\Phi(\mathbf{s}) \geq \|\mathbf{s}\|^2 \geq \mathbf{s}_1^2$. By the assumption on \mathbf{s} , there exists m such that

$$\mathbf{s}_m \geq \mathbf{s}_1/2 \geq \mathbf{s}_{m+1}.$$

Then

$$\|\mathbf{s}\|^2 \geq \frac{m}{4} \mathbf{s}_1^2. \quad (14)$$

By convexity of f ,

$$\sum_{j=1}^m f(-2 + \mathbf{s}_j - \mathbf{s}_{j+1}) \geq \sum_{j=1}^m f\left(-2 + \frac{\mathbf{s}_1 - \mathbf{s}_{m+1}}{m}\right) \geq m f\left(-2 + \frac{\mathbf{s}_1}{2m}\right). \quad (15)$$

We claim (15) is monotonically decreasing in m . To see this, note f' is itself positive and monotonically increasing by monotonicity of σ . Then for $x > -2$,

$$f(x) = \int_0^x f'(t) dt \leq x f'(x) \leq (x+2) f'(x) \implies 0 \leq \frac{(x+2)f'(x) - f(x)}{(x+2)^2} = \frac{d}{dx} \frac{f(x)}{x+2}.$$

Setting $x = -2 + \mathbf{s}_1/2m$ and noting x is itself monotonically decreasing in m establishes the claim. m is the index of a player, so we must have $m \leq n$. This immediately implies by monotonicity of (15) that

$$\Phi(\mathbf{s}) \geq n f\left(-2 + \frac{\mathbf{s}_1}{2n}\right).$$

Combining (14) and (15) yields

$$\begin{aligned}\Phi(\mathbf{s}) &\geq \inf_m \left(\frac{m}{4} \mathbf{s}_1^2 + mf\left(-2 + \frac{\mathbf{s}_1}{2m}\right) \right) \\ &\geq \inf_m \max\left(\frac{m}{4} \mathbf{s}_1^2, mf\left(-2 + \frac{\mathbf{s}_1}{2m}\right)\right).\end{aligned}$$

The minimum of the maximum of two functions occurs when they intersect, which in this case is guaranteed to happen exactly once since $f(-2 + \mathbf{s}_1/2m)$ is monotone in m and its range contains $[0, \infty) \ni \mathbf{s}_1^2/4$. Therefore the minimizing m is

$$m = \frac{\mathbf{s}_1/2}{f^{-1}(\mathbf{s}_1^2/4) + 2}.$$

The last result of the lemma follows immediately by plugging that m into (14). \square

Assembling the above lemmas together results in Theorem 4.5.

Theorem 4.5. *Let r be the highest rating achieved by a group of any number of players who play a total of k games. Suppose the pot function σ satisfies*

$$C_1 = \sup_z z \sigma(2 - z) < \infty.$$

Then

$$r \leq 2n \cdot (f^{-1}(C_2 \cdot k/n) + 2).$$

and

$$r \leq 2 \cdot C_2^{1/3} \cdot k^{1/3} \cdot (f^{-1}(C_2/4 \cdot k) + 2)^{1/3}$$

for $C_2 = 2 + 2C_1 + f(-2)$ where f is defined in (4).

Proof. We may take $n \leq k - 1$ players without loss of generality by simply ignoring players who aren't connected to the player achieving the highest rating via some sequence of games. Let $\mathbf{r} : [0, c] \rightarrow \mathbb{R}^n$ be the valid path corresponding to the sequence of games achieving rating $r = \max \mathbf{r}(c)$. Perform all necessary reflections to make it ordered. Lemma 4.1 gives

$$\text{len}(\mathbf{r}) \leq k. \tag{16}$$

Lemma 4.2 constructs $\tilde{\mathbf{r}}$ without any upsets satisfying

$$r = \max \mathbf{r}(c) \leq \max \tilde{\mathbf{r}}(\tilde{c}) \tag{17}$$

and

$$\text{len}(\tilde{\mathbf{r}}) \leq \text{len}(\mathbf{r}). \tag{18}$$

Since $\tilde{\mathbf{r}}$ is ordered and upset-free, Lemma 4.3 implies

$$\Phi(\tilde{\mathbf{r}}(\tilde{c})) - \Phi(\tilde{\mathbf{r}}(0)) \leq (2 + 2C_1) \text{len}(\tilde{\mathbf{r}}).$$

Note $\Phi(\tilde{\mathbf{r}}(0)) = \Phi(0) = (n - 1)f(-2)$ so

$$\Phi(\tilde{\mathbf{r}}(\tilde{c})) \leq (2 + 2C_1) \text{len}(\tilde{\mathbf{r}}) + (n - 1)f(-2). \tag{19}$$

Set $C_2 = 2 + 2C_1 + f(-2)$ and assemble the chain of inequalities:

$$\begin{aligned}\Phi(\tilde{\mathbf{r}}(\tilde{c})) &\stackrel{(19)}{\leq} (2 + 2C_1) \text{len}(\tilde{\mathbf{r}}) + (n - 1)f(-2) \\ &\stackrel{(18)}{\leq} (2 + 2C_1) \text{len}(\mathbf{r}) + (n - 1)f(-2) \\ &\stackrel{(16)}{\leq} (2 + 2C_1)k + (n - 1)f(-2) \\ &\leq (2 + 2C_1 + f(-2)) \cdot k \\ &= C_2 \cdot k\end{aligned} \tag{20}$$

The three lower bounds on Φ provided by Lemma 4.4 are each used in different ways. First, the n -dependent bound implies

$$nf\left(-2 + \frac{r}{2n}\right) \stackrel{(17)}{\leq} nf\left(-2 + \frac{\max \tilde{\mathbf{r}}(\tilde{c})}{2n}\right) \stackrel{4.4}{\leq} \Phi(\tilde{\mathbf{r}}(\tilde{c})) \stackrel{(20)}{\leq} C_2 \cdot k.$$

Since r appears only once in the equation, we can simply rearrange to solve for r . In particular,

$$r \leq 2n \left(f^{-1} \left(C_2 \cdot \frac{k}{n} \right) + 2 \right).$$

giving the first result of the theorem. The second n -independent bound gives

$$\frac{r^3/8}{f^{-1}(r^2/4) + 2} \stackrel{(17)}{\leq} \frac{(\max \tilde{\mathbf{r}}(\tilde{c}))^3/8}{f^{-1}((\max \tilde{\mathbf{r}}(\tilde{c}))^2/4) + 2} \stackrel{4.4}{\leq} \Phi(\tilde{\mathbf{r}}(\tilde{c})) \stackrel{(20)}{\leq} C_2 \cdot k. \quad (21)$$

This bounds k in terms of a function of r . The last application of Lemma 4.4 converts this into a bound on r in terms of k .

$$r^2 \stackrel{(17)}{\leq} (\max \tilde{\mathbf{r}}(\tilde{c}))^2 \stackrel{4.4}{\leq} \Phi(\tilde{\mathbf{r}}(\tilde{c})) \stackrel{(20)}{\leq} C_2 \cdot k.$$

This can be plugged back into the left-hand side of (21),

$$C_2 \cdot k \geq \frac{r^3/8}{f^{-1}(r^2/4) + 2} \geq \frac{r^3/8}{f^{-1}(C_2/4 \cdot k) + 2}.$$

Rearranging gives

$$r^3 \leq 8 \cdot C_2 \cdot k \cdot (f^{-1}(C_2/4 \cdot k) + 2).$$

Taking cube roots establishes the final result. \square

Remark. The probabilistic interpretation (3) of the Elo system described in the introduction lends itself naturally to the requirement that $z\sigma(2-z)$ be bounded. In particular, take $\sigma(z) = Pr(\eta < x) + \frac{1}{2}Pr(\eta = x)$ to be the symmetrized cumulative distribution function of a symmetric random variable η **with finite expectation**. Note

$$\sup_z z\sigma(2-z) = \sup_z (2+z)\sigma(-z) \leq 2 + \sup_z z\sigma(-z).$$

It's clear that the supremum on the right will be achieved for $z \geq 0$. Since η has finite expectation, we may apply Markov's inequality to obtain

$$z\sigma(-z) \leq zPr(\eta \leq -z) = \frac{1}{2} \cdot z \cdot Pr(|\eta| \geq z) \leq \frac{1}{2} \mathbb{E}|\eta| < \infty.$$

This shows $\mathbb{E}|\eta| < \infty$ is sufficient. However it isn't strictly necessary; for instance, $z\sigma(2-z)$ just barely does not diverge for a Cauchy random variable. However, one *does* need finite $(1-\varepsilon)$ th moment: suppose $\sigma(-z) \leq c/z$. Then

$$\mathbb{E}|\eta|^{1-\varepsilon} = \int Pr(|\eta|^{1-\varepsilon} \geq t) dt = 2 \int \sigma(-t^{1/(1-\varepsilon)}) dt \leq 2 \int t^{-1/(1-\varepsilon)} dt < \infty$$

for each $\varepsilon \in (0, 1)$.

5 Discussion

5.1 Remaining questions

Can one close the gap between the upper and lower bounds in the $n = \omega(1)$ regime? This could occur by finding a better strategy than the ones in Section 3, for instance by slowly increasing δ as games are played,

or by tightening the analysis in Theorem 4.5. It's possible that for heavy-tailed σ the lower bound is too loose, but for light-tail σ the upper bound is. Why is there a jump in the upper bound from $f(k)^{1/3}$ to $f(k/n)$ at $n = k^{1/3}$? Can one find a bound that smoothly crosses the phase transition? The business about having to possibly reflect the set of initial ratings appears very strongly to be an artifact of the analysis. Can it be removed?

The highest Elo problem also bears a passing resemblance to the Toda lattice: players are particles whose position on the real line is given by their ratings; they exhibit a repulsive force when a higher rated player beats a lower rated player and an attractive force if vice versa. Is it possible to use tools from solid state physics to analyze this problem?

Another interesting variant is to constrain the number of *rounds* of games. In a tournament, many games will be happening parallel. Each round consists of any number of games, subject only to the constraint that each player only participates in at most one game each round. **Given n players and k rounds, what's the highest someone may be rated after all games have finished?**

5.2 Connection to maximum overhang

One may notice the jump from $\log k$ to $k^{1/3}$ as one increases n for $\sigma(-z) = \frac{1}{1+e^{-z}}$. This may be reminiscent of the maximum overhang problem. In that problem, one places k unit-length bricks on top of each other at the edge of a table, attempting to achieve the largest overhang possible. The classic solution if one is restricted to placing only a single brick at each height (the 'single-wide' setting) achieves an overhang of $\Theta(\log k)$ units. However, if one is allowed to place as many bricks at each height as one likes (the 'multi-wide' setting), the optimal solution achieves an overhang of $\Theta(k^{1/3})$ [5]. This connection does not appear to be a coincidence! The proof of optimality in the maximum overhang problem uses a reduction to a 'mass movement problem.' But in fact, the highest Elo problem for the particular $\sigma(z) = \max(0, \min(1, z/2 + 1/2))$ can be reduced to nearly the same problem!

The mass movement problem described by [5] is as follows: consider some finite number of piles of mass placed on the real line. A valid 'move' takes some unit interval and rearranges the mass within that interval so that the center of mass is unchanged. Negative mass is allowed. Formally, a 'signed distribution' μ is a finite linear combination of dirac delta functions δ . A valid 'move' replaces μ with $\mu + \nu$ where ν is itself a signed distribution whose support is contained in some unit interval and $\int \nu = 0$. A sequence of moves $\nu_1 \cdots, \nu_k$ corresponds naturally to a sequence of signed distributions μ_0, \cdots, μ_k . They place an additional 'weight-constraint' on allowed sequences. In particular, the number of j for which ν_j is allowed to have support to the right of any threshold T is at most $\max_{0 \leq i \leq k} \int_T^\infty \mu_i$. The goal is to use at most k moves to move the distribution $k \cdot \delta$ to a distribution with a unit of mass as far right as possible.

To reduce highest Elo to mass movement, let $r_p(j)$ be the rating of player p after j games and set

$$\mu_j(x) = \sum_p \delta(x - r_p(j)).$$

Then consider a game where p beats q . Set $z = r_p - r_q$. By the selection of $\sigma(z) = \max(0, \min(z/2 + 1/2))$, this only produces a change in ratings if $z \leq 1$. Then μ changes by performing the following move:

$$\nu(x) = \delta(x - r_p - \sigma(-z)) + \delta(x - r_q + \sigma(-z)) - \delta(x - r_p) - \delta(x - r_q).$$

Suppose $z \geq -1$. Then the support of ν is contained in $[r_q - \sigma(-z), r_p + \sigma(-z)]$, which is an interval of length $z + 2\sigma(-z) = 1$ as required. Furthermore, ν only has support above T when there's a player above T Elo guaranteeing the weight-constraint. The objective is starting with $n \cdot \delta$, produce via k moves a distribution with a unit of mass (i.e. at least one player) as far right as possible (i.e. with the highest rating possible). Recall that for the purposes of an upper bound, we make take $n = k$, establishing the correct objective.

There is, however, one big discrepancy: the case of $z < -1$. In particular, in the highest Elo problem it is possible for a player rated -10 to beat a player rated 10 , garnering a full point of Elo. This does not correspond to a valid move, since the support of ν would be $\{-10, -9, 9, 10\}$. Intuitively, such moves should

not ultimately help push mass far away, and indeed in Lemma 4.2 we show for a different relaxation of highest Elo that they don't help.

A second discrepancy is that mass movement places no restriction on the amount of mass moved in each move, whereas for highest Elo we always have $\|\nu\|_1 = 2$. This would seem to imply an upper bound of the ' k rounds' variant mentioned at the end of Section 5.1. Can that argument be made to work? Conversely, is there a variant or generalization of maximum overhang that can be analyzed using the proof of Theorem 4.5?

References

- [1] Arpad E. Elo. The proposed USCF rating system: Its development, theory, and applications. *Chess Life*, 23(8), August 1967.
- [2] Arpad E. Elo. *The Rating of Chess Players, Past and Present*. Arco Publishing, second edition, 1978.
- [3] FIDE. FIDE Ratings and Statistics. <https://ratings.fide.com/>.
- [4] FIDE. FIDE Handbook. <https://handbook.fide.com/>, 2017.
- [5] Mike Paterson, Yuval Peres, Mikkel Thorup, Peter Winkler, and Uri Zwick. Maximum overhang. *The American Mathematical Monthly*, 116(9):763–787, 2009.
- [6] Jeff Sonas. The Elo rating system – correcting the expectancy tables. *ChessBase*, March 2011.