Becoming the World's Highest Rated Chess Player

Rikhav Shah

January 2022

Abstract

The Elo rating system measures the approximate skill of each competitor in a game or sport. As a competitor competes, their rating changes based on their wins and losses. 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. 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? The people you gathered aren't very interested in the game, and will only play if you pay them to. We resolve this question for $n = 2$ and $n = \Omega(k^{1/3})$ (up to a log factor) for the standard Elo system, and provide close upper and lower bounds for a wider class of rating systems that resemble Elo.

1 Introduction

The Elo Rating system was proposed by Arpad Elo in the mid 20th century to accurately estimate the relative skill of chess players [1]. 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 asked: 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 description of the system. Each player is given some 'rating' value (measured in 'points'), 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. Rather than decide how many points to stake themselves, the Elo system posits that the amount a player puts in should be proportional to the estimated probability that player will win the game, given the difference in ratings between them and their opponent. Specifically, the system establishes a 'pot function' σ which requires players A and B place $\sigma(r_A - r_B)$ and $\sigma(r_B - r_A)$ points into the pot respectively, where r_A, r_B are their ratings. The usual pot function used is $\sigma(z) = \frac{1}{1+e^{-z}}$ (up to scaling), but our results hold for a much broader class. Real-world uses of Elo typically avoid fractional and negative rating points 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 size of the pot may also vary depending on their ratings or on how many games the player has played before. For simplicity, we ignore these complications by allowing for both fractional and negative points, and fixing a pot size of 1 point. Since the dynamics depend only on the difference between ratings, we can take the initial rating to be 0 without loss of generality.

Pot function	$n=2$	$n=\Omega(k^{1/3})$	$n=\infty$
$\sigma(z)$	$\frac{1}{2}f^{-1}(2k)\pm O(1)$	$\Omega(k^{1/3})$	$O(k^{1/3}g^{-1}(ck)^{1/3})$*
$\frac{1}{1+e^{-z}}$	$\frac{1}{2}\log k \pm O(1)$	$\Omega(k^{1/3})$	$O(k^{1/3}\log^{1/3}(k))$
$\frac{1}{2}\frac{z}{(1+\lvert z\rvert^p)^{1/p}}+\frac{1}{2}$	$\frac{1}{2}\left(1+\frac{1}{p}\right)^{\frac{1}{p+1}}k^{\frac{1}{1+p}}\pm O(1)$	$\Omega(k^{1/3})$	$O\left(k^{1/3+\frac{2}{3}\frac{1}{3p+1}}\right)$
$\frac{1}{2}\mathrm{erf}(z/\sqrt{2})+\frac{1}{2}$	$\frac{1}{2}\sqrt{\log k}\pm O(1)$	$\Omega(k^{1/3})$	$O(k^{1/3}\log^{1/6}(k))$

Table 1: The values of the largest rating achievable with n players and k games for some selected pot functions. In the third row, the parameter p is constrained to $p \geq 1$. (*) The general case for $n=\infty$ holds when $g(x)=\Omega(x^2)$, and otherwise is the inverse of the function $x \mapsto x^3/g^{-1}(x)$. The definitions of f and g are given by (1). The case of $n=2$ is covered by Theorem 1, the $n=\Omega(k^{1/3})$ case by Theorem 2, and the $n=\infty$ case by Corollary 7.

In the case of $\sigma(z)=\frac{1}{1+e^{-z}}$, we show the maximum rating is $\log(k)/2 \pm O(1)$ when $n=2$, and $\tilde{\Theta}(k^{1/3})$ when $n=\Omega(k^{1/3})$. Results for other natural choices of σ are contained in Table 1. One can intuitively think of $\sigma(r_A - r_B)$ as the estimated probability A will beat B. The outcome of a game between A and B is modeled as a random variable $r_A - r_B + \eta$ for symmetric η, where A beats B if and only if $r_A - r_B + \eta > 0$. Then, σ is simply the cumulative distribution function of η. Throughout the paper, we make four assumptions of the pot function, which correspond to mild regularity assumptions on η.

1. σ is positive and monotonically increasing everywhere.
2. $\sigma(z)+\sigma(-z)=1$ for all z.
3. $\sigma(2-2z) < z^{-1}$ for $z > 0$.
4. $\sup_{z\geq 0} \frac{\sigma(-z)}{\sigma(-z-2\sigma(-z))} < \infty$.

The fourth condition is the strangest seeming, but for a number of pot functions, including all ones listed in Table 1, it appears the supremum is achieved for $z=0$, making the condition usually unnecessary. Associated with each σ are the following functions which roughly capture how quickly the left tail of σ tends toward 0.

$$f(x)=\int_0^x \frac{1}{\sigma(-\tau)}\,d\tau, \qquad g(x)=\int_0^x \tau f''(\tau)\,d\tau \tag{1}$$

where f'' denotes the second derivative of f. Note that f is convex and monotonically increasing, and that and g is monotonically increasing for $x \geq 0$.

2 Case of $n=2$

When there are only two players, we can analyze what happens when one player repeatedly beats a second. As their ratings grow further apart, the fewer rating points the first player earns from the second each time. However, the amount is positive no matter the value of their ratings, so the first player's rating indeed tends to infinity as k grows. The question is at what rate. One can imagine that the faster the left tail of σ decays to 0, 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 (1).

Theorem 1. *Say two players start at 0 rating and the first beats the second k times consecutively. Let $r(k)$ be the first player's rating after the k games. Then*

$$\left|\frac{1}{2}f^{-1}(2k) - r(k)\right| \leq 3$$

where f is given by (1).

Proof. The ratings of the two players sum to zero, so we have a simple recurrence

$$r(t+1) = r(t) + \sigma(-2r(t)).$$

By monotonicity of σ, when $r(t) \in [a-1, a)$, we have

$$\sigma(-2a) \leq r(t+1) - r(t) \leq \sigma(-2(a-1)),$$

so there are at least $\lfloor 1/\sigma(-2(a-1)) \rfloor$ and at most $\lceil 1/\sigma(-2a) \rceil$ values of t for which $r(t) \in [a-1, a)$ holds. To obtain bounds on the number of values of t for which $r(t) < a$, we sum over the intervals from 0 to a and use the property of σ that $a \leq 1/\sigma(2-2a) \leq 1/\sigma(-2a-2)$ for positive a,

$$|\{t \,|\, r(t) < a\}| \leq \sum_{j=1}^{a} \left\lceil \frac{1}{\sigma(-2j)} \right\rceil \leq \ a + \sum_{j=1}^{a} \frac{1}{\sigma(-2j)} < \sum_{j=1}^{a+1} \frac{1}{\sigma(-2j)} \leq f(2a+4)/2,$$

$$|\{t \,|\, r(t) < a\}| \geq \sum_{j=0}^{a-1} \left\lfloor \frac{1}{\sigma(-2j)} \right\rfloor \geq -a + \sum_{j=0}^{a-1} \frac{1}{\sigma(-2j)} > \sum_{j=0}^{a-2} \frac{1}{\sigma(-2j)} \geq f(2a-4)/2.$$

Since $r(t)$ is monotone in t, the set $\{t \,|\, r(t) < a\}$ is simply all the natural numbers up to $\max\{t \,|\, r(t) < a\}$. Then by definition,

$$r(\max\{t \,|\, r(t) < a\}) < a \leq r(\max\{t \,|\, r(t) < a\} + 1).$$

The inequalities we just established give

$$r(\lceil f(2a-4)/2 \rceil) < a \leq r(\lceil f(2a+4)/2 \rceil)$$

or equivalently

$$a - 2 \leq r(\lceil f(2a)/2 \rceil) < a + 2.$$

Let k be any integer in the interval $[f(2a)/2, f(2(a+1))/2]$. Then

$$a - 2 \leq r(k) < a + 3$$

and additionally, $f^{-1}(2k)/2$ is in the interval $[a, a+1]$, so we conclude

$$\frac{1}{2}f^{-1}(2k) - 3 \leq r(k) < \frac{1}{2}f^{-1}(2k) + 3.$$

□

3 Case of $n = \Omega(k^{1/3})$

The previous section examined the 'extremal' regime of $n = 2$. The other extremal regime is $n = \infty$, that is, we're allowed to use as many players as we need. This section describes an explicit strategy using $\Theta(k^{1/3})$ players. Surprisingly, the asymptotic rate at which the maximum rating increases as k increases is mostly independent of the particular pot function used. In all cases, we achieve a rate of $\Omega(k^{1/3})$. The proof of this fact first assumes access to an infinite pool of players, and then bounds the number of players who actually need to play any games.

Theorem 2. *For each k sufficiently large, there exists a strategy in which a total of k games are played among some number of players starting at rating 0 such that at least one has rating at least $c_1 k^{1/3}$ after the games. Furthermore, the strategy requires only $c_2 k^{1/3}$ players. The constants c_1, c_2 depend only on the pot function.*

Proof. If $\lim_{z\to\infty} \sigma(-z)z^2 \not\to 0$, then $f(x)$ defined by (1) is $O(x^3)$ giving $r(k) = \Omega(k^{1/3})$ with just two players by Theorem 1. In the remaining cases, there exists positive z achieving the maximum of $\sigma(-z)z^2$. For the best asymptotic rate, put $A = \arg\max_{z\geq 0} \sigma(-z)z^2$, though the analysis is the same for any $A > 0$. Arrange an infinite number of chess players in a line, numbered 1 through infinity. Let r_i denote the rating of the ith player, initialized to 0 for all i. We describe the strategy explicitly: repeat the following two steps k times,

1. Find the smallest i such that $r_i < r_{i+1} + A$.
2. Have player i beat $i + 1$.

Let $n_k + 1$ be the largest index of any player who played at least one game. Assume k is large enough for $n_k > 1$. Note that throughout the procedure r_1 only increases, so if we give it a lower bound at any intermediate point in time, that will serve as a lower bound on its final value. The time we pick is just before the $(n_k + 1)$th player plays their first game. Since n_k is the smallest index found in Step 1, at this point in time we have

$$A \leq r_i - r_{i+1} \quad \forall i \in [n_k - 1]. \tag{2}$$

Define the potential function

$$\phi_k(\mathbf{r}) = \sum_{i=1}^{n_k} (2n_k - 2i + 1) r_i.$$

Observe that when x Elo is transferred from player $i+1$ to player i, ϕ_k increases by $2x$. Also recall that that transfer only occurs when $r_{i+1} - r_i > -A$, so the amount of Elo transferred is at least $\sigma(-A)$. Since $\phi_k(0) = 0$, this gives a lower bound of

$$2\sigma(-A)k \leq \phi_k(\mathbf{r}). \tag{3}$$

We can also rewrite ϕ_k as a telescoping sum using the fact that $(2n_k - 2i + 1) = (n_k - i + 1)^2 - (n_k - i)^2$,

$$\phi_k(\mathbf{r}) = n_k^2 r_1 + \sum_{i=1}^{n_k - 1} (n_k - i)^2 (r_{i+1} - r_i).$$

Now apply (2) and rearrange to obtain a lower bound on r_1,

$$\frac{1}{n_k^2}\left(\phi(\mathbf{r}) + A\sum_{i=1}^{n_k-1}(n_k-i)^2\right) \leq r_1.$$

Now use (3) and $n_k^3/3 - n_k^2 \leq \sum_{i=1}^{n_k-1}(n_k-i)^2$ then distribute,

$$2\sigma(-A)\frac{k}{n_k^2} + \frac{A}{3}n_k - A \leq r_1. \tag{4}$$

The left-hand side is a convex function of n_k, and is minimized for $n_k = (12\sigma(-A)/A)^{1/3}k^{1/3}$. Set $c_1 = (\sigma(-A)A^2)^{1/3}$. Plugging this value in yields

$$r_1 \geq \frac{2^{-1/3}+2^{2/3}}{9^{1/3}}c_1k^{1/3} - A > 1.14c_1k^{1/3} - A \gtrsim c_1k^{1/3}. \tag{5}$$

To bound the number of players required, observe the invariant that the sum of all ratings is 0. Note $r_i = 0$ for $i > n_k$, so we may ignore those players. We can rearrange the sum of all the ratings as a telescoping sum of differences,

$$0 = \sum_{i=1}^{n_k} r_i = n_k r_1 + \sum_{i=1}^{n_k-1}(n_k-i)(r_{i+1}-r_i).$$

Applying (2), rearranging, then using $\sum_{i=1}^{n_k-1}(n_k-i) = n_k(n_k-1)/2$ gives

$$\frac{A}{2}(n_k-1) \leq r_1. \tag{6}$$

Lastly, note that we may end the strategy as soon as r_1 exceeds $1.14c_1k^{1/3} - A$. Since r_1 increases by at most 1 each time, this means $r_1 \leq 1.14c_1k^{1/3} - A + 1$ at that stopping point, and the number of players who play any games is upper bounded by n_k, and by (6),

$$n_k \leq 1.14\frac{2}{A}c_1k^{1/3} - 1 \lesssim c_2k^{1/3}$$

giving the final result. □

4 Optimality ($n = \infty$)

In this section, we wish to show the algorithm presented in Theorem 2 is optimal in the sense that with k games and an *infinite* number of players, one cannot achieve a higher rating than $O(k^{1/3})$. We will argue the dual statement: that achieving a rating of R requires at least (roughly) $\Omega(R^3)$ games to be played. To do this, we will provide a lower bound on lengths of paths in a particular directed graph.

Consider the space of all possible vectors of ratings of n players as a subset of $\mathbb{R}^n$. We can construct a directed graph by connecting each vector of ratings to the possible resulting vectors after one game is played. Specifically, each vertex $\mathbf{r}$ of the graph is connected to $\mathbf{r}+(e_i-e_j)\sigma(r_j-r_i)$

for each i, j, corresponding to the resulting vector of ratings if i beats j. Let $T_R = \bigcup_i \{\mathbf{r} \mid r_i \geq R\}$ be the set of all vectors with at least one coordinate at least R. Then we wish to show that the length of the shortest path from the origin to any point in T_R is some quickly growing function of R. Now, adding edges and vertices can only decrease this length. We extend the vertex set to be all of $\mathbb{R}^n$ and add edges from $\mathbf{r}$ to

$$\mathbf{r} + (e_i - e_j)\sigma(r_j - r_i)t$$

for each $i, j \in [n], t \in (0, 1]$, where the weight of the edge is t. Note that the case of $t = 1$ recovers the edges of the original graph. The intuition is that t measures the size of the pot the two players are vying for. In the standard setup, t is always 1, but now we consider a 'half'-game to be played if the players make the total pot of size 1/2, and similarly a t^{th} of a game to be played when the pot has size t. Also connect $\mathbf{r}$ to $P\mathbf{r}$ with an edge of weight 0 for every permutation matrix P. The intuition is that these edges correspond to the players standing up and rearranging themselves. Call the resulting graph G. Let $\text{len}(p)$ denote the weighted length of a path p in G, that is, the sum of the weights of the edges in the path. Note that the edges in this graph corresponding to games can be described as $(v, t) \in \{e_i - e_j \mid i, j \in [n]\} \times (0, 1]$ denoting the direction and magnitude of each move. Specifically, (v, t) denotes the edge $(\mathbf{r}, \mathbf{r} + v\sigma(-\langle \mathbf{r}, v\rangle)t)$ for some $\mathbf{r}$. Let $\pi_{\mathbf{r}}$ denote a permutation which sorts the entries of $\mathbf{r}$ into non-increasing order, i.e. $r_{\pi_{\mathbf{r}}(i)} \geq r_{\pi_{\mathbf{r}}(i+1)}$. Then define

$$\Phi(\mathbf{r}; \pi) = \|\mathbf{r}\|^2 + \sum_{\ell=1}^{n-1} f(r_{\pi(\ell)} - r_{\pi(\ell+1)})$$

for f given by (1). We will frequently consider $\Phi(\mathbf{r}; \pi_{\mathbf{r}})$. Finally, call an edge $(\mathbf{r}, \mathbf{r} + v\sigma(-\langle \mathbf{r}, v\rangle)t)$ an 'upset' if $\langle \mathbf{r}, v\rangle < 0$ (or equivalently $r_i < r_j$ for $v = e_i - e_j$), i.e. the lower rated player wins the game. An upset-free path is a path which does not contain any upsets.

At a high level, the proof first argues that any path from the origin to T_R can be converted to an upset-free path. Then, the length of an upset-free path is lower bounded in terms of some expression of Φ, which in turn is lower bounded by some fast growing function of R. These three steps are separated into the three lemmas 3, 4, 5, when then conclude in Theorem 6.

Lemma 3. *Let p be any path in G from the origin to T_R. Then there exists a upset-free path p' from the origin to T_R such that $\text{len}(p') \leq C_1 \text{len}(p)$. The constant C_1 may depend on σ.*

Proof. We construct p' from p by making several passes over the edges in p and making modifications along the way. First, we remove all of a certain kind of upset. Then, we show the remaining upsets can be moved to the end of the path. Lastly, we show upsets at the end of the path can be completely removed.

Step 1: Removing a certain kind of upset. Let i be the index of the winning player and j the losing. Let r_ℓ and r'_ℓ respectively denote the rating of the ℓth player before and after the game. The upsets we first remove are ones for which $r'_i > r'_j$, that is, the winning player gains enough Elo to surpass the losing player. There are a few sub-cases to consider.

1. $r'_i = r_j$ and $r'_j = r_i$. That is, i gains enough Elo to match j's old rating, and vice versa. We can simply replace this edge with a cost 0 transposition of i and j.

2. $r'_i > r'_j$ and $r'_i - r'_j < r_j - r_i$. This case corresponds to the game transferring enough Elo for i to surpass j in rating, but not enough to surpass j's old rating. We replace this edge with

pair of edges: the first is $(\mathbf{r}, \mathbf{r} + (e_i - e_j)\sigma(r_j - r_i)\gamma)$ that is, the same pair of players play, just a γ^{th} of a game instead of a t^{th}. The second edge is a transposition of i and j. We want to show the end of these two two edges is still $\mathbf{r}'$, and further that $0 < \gamma \leq t$. The end point condition requires

$$r_i + \sigma(r_j - r_i)t = r_j - \sigma(r_j - r_i)\gamma,$$
$$r_j - \sigma(r_j - r_i)t = r_i + \sigma(r_j - r_i)\gamma.$$

Taking the difference and replacing $z = r_j - r_i$ gives

$$z - 2\sigma(z)t = -z + 2\sigma(z)\gamma$$

so $\gamma = \frac{z}{\sigma(z)} - t$. The condition that i surpasses j gives $0 > r'_j - r'_i = z - 2\sigma(z)t$. Rearranging gives $\gamma < t$. The condition that i does not surpass j's old rating gives $2\sigma(z)t - z = r'_i - r'_j < r_j - r_i = z$. Rearranging gives $\gamma > 0$.

3. $r'_i > r'_j$ and $r'_i - r'_j > r_j - r_i$. In this case, i surpasses j's old rating. Then replace the edge with a non-upset between the two players followed by a transposition of i and j. The first inserted edge is $(\mathbf{r}, \mathbf{r} + (e_j - e_i)\sigma(r_i - r_j)\gamma)$. The condition that the endpoints match is

$$r_i + \sigma(r_j - r_i)t = r_j + \sigma(r_i - r_j)\gamma,$$
$$r_j - \sigma(r_j - r_i)t = r_i - \sigma(r_i - r_j)\gamma.$$

Taking differences and putting $z = r_j - r_i$ gives

$$z - 2\sigma(z)t = -z - 2\sigma(-z)\gamma$$

so $\gamma = t - \frac{z}{\sigma(z)}$. The condition that i surpasses j's old rating gives $\gamma > 0$. Then $z > 0$ gives $\gamma < t$.

In each of those cases, the length of the path only decreases.

Step 2: Moving remaining upsets to the end. We now introduce a new notion of length. For a non-upset (v, t), let it's new cost, called the 'continuous cost', be

$$f(z + 2\sigma(-z)t) - f(z) = \int_z^{z+2\sigma(-z)t} \frac{d\tau}{\sigma(-\tau)}$$

where $z = \langle \mathbf{r}, v\rangle$. First observe that the continuous cost is 0 when $t = 0$. By the fourth property of σ, the continuous cost is an increasing function of t with derivative bounded independently of z. As a consequence, the continuous cost is upper bounded by a constant times the real cost t. Additionally, sub-dividing (v, t) into edges $(v, \gamma_1), (v, \gamma_2)$ has no effect on the continuous cost. Let the continuous length of a path be the sum of the continuous costs of the non-upsets, ignoring entirely the upsets. Our next modifications to the path will reduce its continuous cost. Consider two consecutive edges $(u, t_u), (v, t_v)$ for an upset u and non-upset v. If no such pair exists, proceed to Step 3. Denote the corresponding sequence of three vertices by

$$\begin{aligned}
&\mathbf{r} \\
\mathbf{r}' &:= \mathbf{r} + u\sigma(-\langle \mathbf{r}, u\rangle)t_u \\
\mathbf{r}'' &:= \mathbf{r} + u\sigma(-\langle \mathbf{r}, u\rangle)t_u + v\sigma\left(-\langle \mathbf{r}', v\rangle\right)t_v \\
&= \mathbf{r} + u\sigma(-\langle \mathbf{r}, u\rangle)t_u + v\sigma\left(-\langle \mathbf{r}, v\rangle - \langle u, v\rangle\sigma(-\langle \mathbf{r}, u\rangle)t_u\right)t_v.
\end{aligned}$$

The continuous cost is

$$\int_{\langle \mathbf{r}', v\rangle}^{\langle \mathbf{r}', v\rangle + 2\sigma\left(-\langle \mathbf{r}', v\rangle\right) t_v} \frac{d\tau}{\sigma(-\tau)}. \tag{7}$$

The cases we consider correspond to the possible values of $\langle u, v\rangle$. The case of $\langle u, v\rangle = 2$ was removed by Step 1—an upset will not cause the winning player to surpass the losing player, so $u \neq v$. When $\langle u, v\rangle = 0$, the last vertex is $\mathbf{r}'' = \mathbf{r} + u\sigma(-\langle \mathbf{r}, u\rangle)t_u + v\sigma(-\langle \mathbf{r}, v\rangle)t_v$, so we can simply swap the order of the edges with no effect on cost. When $\langle u, v\rangle = 1$, replace the edges with $(v, \gamma_v), (u, \gamma_u)$ where

$$\gamma_v = \frac{\sigma(-\langle \mathbf{r}', v\rangle)}{\sigma(-\langle \mathbf{r}, v\rangle)} t_v, \qquad \gamma_u = \frac{\sigma(-\langle \mathbf{r}, u\rangle)}{\sigma\left(-\langle \mathbf{r}, u\rangle - \sigma\left(-\langle \mathbf{r}, v\rangle\right)\gamma_v\right)} t_u.$$

Then the new end point is

$$\mathbf{r} + v\sigma\left(-\langle \mathbf{r}, v\rangle\right)\gamma_v + u\sigma\left(-\langle \mathbf{r}, u\rangle - \sigma\left(-\langle \mathbf{r}, v\rangle\right)\gamma_v\right)\gamma_u$$

which indeed equals the old end point $\mathbf{r}''$ for our choice of γ_u, γ_v. The new continuous cost is

$$\int_{\langle \mathbf{r}, v\rangle}^{\langle \mathbf{r}, v\rangle + 2\sigma(-\langle \mathbf{r}, v\rangle)\gamma_v} \frac{d\tau}{\sigma(-\tau)} = \int_{\langle \mathbf{r}, v\rangle}^{\langle \mathbf{r}, v\rangle + 2\sigma(-\langle \mathbf{r}', v\rangle)t_v} \frac{d\tau}{\sigma(-\tau)}.$$

Comparing the above to (7), observe that the length of the integration path is the same. Since $\langle \mathbf{r}', v\rangle \geq \langle \mathbf{r}, v\rangle$, the new cost is at most the old cost by monotonicity of σ.

Finally, when $\langle u, v\rangle = -1$, there are two sub-cases. The first sub-case is $\sigma\left(-\langle \mathbf{r}, u\rangle\right)t_u \geq \sigma\left(-\langle \mathbf{r}', v\rangle\right)t_v$. Then replace the edges with $(u + v, \gamma_w), (u, \gamma_u)$ where

$$\gamma_w = \frac{\sigma\left(-\langle \mathbf{r}', v\rangle\right)}{\sigma\left(-\langle \mathbf{r}, u + v\rangle\right)} t_v, \qquad \gamma_u = \frac{\sigma\left(-\langle \mathbf{r}, u\rangle\right)t_u - \sigma\left(-\langle \mathbf{r}', v\rangle\right)t_v}{\sigma\left(-\langle \mathbf{r}, u\rangle - \sigma\left(-\langle \mathbf{r}', v\rangle\right)t_v\right)}. \tag{8}$$

The new endpoint is

$$\begin{aligned}
&\mathbf{r} + (u+v)\sigma\left(-\langle \mathbf{r}, u+v\rangle\right)\gamma_w + u\sigma\left(-\langle \mathbf{r} + (u+v)\sigma\left(-\langle \mathbf{r}, u+v\rangle\right)\gamma_w, u\rangle\right)\gamma_u \\
&= \mathbf{r} + (u+v)\sigma\left(-\langle \mathbf{r}, u+v\rangle\right)\gamma_w + u\sigma\left(-\langle \mathbf{r}, u\rangle - \sigma\left(-\langle \mathbf{r}, u+v\rangle\right)\gamma_w\right)\gamma_u \\
&= \mathbf{r} + (u+v)\sigma\left(-\langle \mathbf{r}', v\rangle\right)t_v + u\sigma\left(-\langle \mathbf{r}, u\rangle - \sigma\left(-\langle \mathbf{r}', v\rangle\right)t_v\right)\gamma_u \\
&= \mathbf{r} + v\sigma\left(-\langle \mathbf{r}', v\rangle\right)t_v + u\left(\sigma\left(-\langle \mathbf{r}', v\rangle\right)t_v + \sigma\left(-\langle \mathbf{r}, u\rangle - \sigma\left(-\langle \mathbf{r}', v\rangle\right)t_v\right)\gamma_u\right) \\
&= \mathbf{r} + v\sigma\left(-\langle \mathbf{r}', v\rangle\right)t_v + u\sigma\left(-\langle \mathbf{r}, u\rangle\right)t_u \\
&= \mathbf{r}''
\end{aligned}$$

as required. The new continuous cost is

$$\int_{\langle \mathbf{r}, u+v\rangle}^{\langle \mathbf{r}, u+v\rangle + 2\sigma(-\langle \mathbf{r}, u+v\rangle)\gamma_w} \frac{d\tau}{\sigma(-\tau)} = \int_{\langle \mathbf{r}, u+v\rangle}^{\langle \mathbf{r}, u+v\rangle + 2\sigma(-\langle \mathbf{r}', v\rangle)t_v} \frac{d\tau}{\sigma(-\tau)}.$$

The value of γ_w guarantees that the length of the integration path is the same as in (7). It just remains to show that the starting points satisfy $\langle \mathbf{r}, u+v\rangle \leq \langle \mathbf{r}', v\rangle$ to conclude that the new cost is lower. Since (u, t_u) is an upset, we have $\langle \mathbf{r}, u\rangle < 0$. Since Step 1 removed all upsets that swap the rank of the two players, we also have $\langle \mathbf{r}', u\rangle \leq 0$. Thus

$$0 \geq \langle \mathbf{r}', u\rangle = \langle \mathbf{r}, u\rangle + 2\sigma\left(-\langle \mathbf{r}, u\rangle\right)t_u \geq \langle \mathbf{r}, u\rangle + \sigma\left(-\langle \mathbf{r}, u\rangle\right)t_u.$$

Then adding $\langle \mathbf{r}', v\rangle = \langle \mathbf{r}, v\rangle - \sigma\left(-\langle \mathbf{r}, u\rangle\right) t_u$ to each side of the inequality gives the desired result.

The second sub-case is $\sigma\left(-\langle \mathbf{r}, u\rangle\right) t_u < \sigma\left(-\langle \mathbf{r}', v\rangle\right) t_v$. Then simply sub-divide the second edge (v, t_v) into $(v, \gamma_v), (v, \gamma_v')$ for

$$\gamma_v = \frac{\sigma\left(-\langle \mathbf{r}, u\rangle\right)}{\sigma\left(-\langle \mathbf{r}', v\rangle\right)} t_u$$

and γ_v' whatever is required to reach the same endpoint. Then the edges $(u, t_u), (v, \gamma_v)$ fall into the first sub-case. However, note that we actually have equality $\sigma\left(-\langle \mathbf{r}, u\rangle\right) t_u = \sigma\left(-\langle \mathbf{r}', v\rangle\right) t_v$, which means the γ_u stated in (8) will be 0. In other words, the upset is removed entirely. The new two edges inserted are thus $(u+v, \gamma_w), (v, \gamma_v')$.

Step 3: Removing ending upsets. We can now assume the last edge is an upset. However, since we removed upsets in which the winning player surpasses the losing player, an upset can only decrease the largest entry in $\mathbf{r}$. So, simply removing the last edge preserves the membership of the final endpoint in T_R. By repeating this process many times, we completely eliminate all upsets, and so the total real cost is upper bounded by the continuous cost, which in turn is at most a constant multiple of the original real cost, as desired. □

Lemma 4. *Let p be an upset-free path from the origin to $\mathbf{r}$. Then*

$$\Phi(\mathbf{r}; \pi_{\mathbf{r}}) \leq \left(6 + \sup_{z \geq 0} \frac{4\sigma(-z)}{\sigma(-z - 2\sigma(-z))} + \frac{1}{\sigma(-1)}\right) \operatorname{len}(p).$$

The assumptions on σ listed in Section 1 ensure that the right-hand side is $C_2 \operatorname{len}(p)$ for some constant C_2 that depends only on σ.

Proof. Note $\Phi(0, \pi) = 0$. Also note $\Phi(\mathbf{r}, \pi_{\mathbf{r}})$ is unaffected by permuting the entries of $\mathbf{r}$. Thus we just need to bound the increase in Φ resulting from edges that correspond to games. Let $(\mathbf{r}, \mathbf{r}' = \mathbf{r} + v\sigma(-\langle \mathbf{r}, v\rangle)t)$ be an edge found in the path. We will show that

$$\Phi(\mathbf{r}'; \pi_{\mathbf{r}'}) - \Phi(\mathbf{r}; \pi_{\mathbf{r}}) \leq \left(6 + \sup_{z \geq 0} \frac{4\sigma(-z)}{\sigma(-z - 2\sigma(-z))} + \frac{1}{\sigma(-1)}\right) t.$$

We first obtain an upper bound on $\Phi(\mathbf{r}'; \pi_{\mathbf{r}}) - \Phi(\mathbf{r}; \pi_{\mathbf{r}})$ (note the first term has $\pi_{\mathbf{r}}$ not $\pi_{\mathbf{r}'}$). Let $z = \langle \mathbf{r}, v\rangle$. Then

$$\begin{aligned}
\|\mathbf{r} + v\sigma(-\langle \mathbf{r}, v\rangle)t\|^2 - \|\mathbf{r}\|^2 &= 2\langle \mathbf{r}, v\rangle \sigma(-\langle \mathbf{r}, v\rangle)t + 2\sigma(-\langle \mathbf{r}, v\rangle)^2 t^2 \\
&= 2z\sigma(-z)t + 2\sigma(-z)^2 t^2 \\
&\leq (2z\sigma(-z) + 2)t \leq 6t.
\end{aligned}$$

For the second term of Φ, we need the assumption that the path is upset-free. Recall that $\pi_{\mathbf{r}}(\ell)$ is the ℓth ranked player, so $\pi_{\mathbf{r}}(i)$ can only beat $\pi_{\mathbf{r}}(j)$ when $i < j$. For concision, put $s_\ell = r_{\pi_{\mathbf{r}}(\ell)}$ and $s'_\ell = r'_{\pi_{\mathbf{r}}(\ell)}$ so s_ℓ is the rating of the ℓth ranked player. For the remainder of the proof, fix indices i, j such that $v = e_{\pi_{\mathbf{r}}(i)} - e_{\pi_{\mathbf{r}}(j)}$, in other words let i denote the rank of the winning player and j the rank of the losing player. The first case is $j = i + 1$. Then since f is monotonically increasing, only one term of the sum in Φ increases:

$$\begin{aligned}\sum_{\ell=1}^{n-1} f(r'_{\pi_{\mathbf{r}}(\ell)} - r'_{\pi_{\mathbf{r}}(\ell+1)}) - \sum_{\ell=1}^{n-1} f(r_{\pi_{\mathbf{r}}(\ell)} - r_{\pi_{\mathbf{r}}(\ell+1)}) &\leq f(s_i - s_j + 2\sigma(s_j - s_i)t) - f(s_i - s_j) \\ &= \int_z^{z+2\sigma(-z)t} \frac{d\tau}{\sigma(-\tau)} \\ &\leq \frac{2\sigma(-z)}{\sigma(-z - 2\sigma(-z))} t \end{aligned} \tag{9}$$

where the last step follows from monotonicity of σ. The remaining cases are $j > i+1$. Then the analysis is similar.

$$\begin{aligned}\sum_{\ell=1}^{n-1} f(r'_{\pi_{\mathbf{r}}(\ell)} - r'_{\pi_{\mathbf{r}}(\ell+1)}) - \sum_{\ell=1}^{n-1} f(r_{\pi_{\mathbf{r}}(\ell)} - r_{\pi_{\mathbf{r}}(\ell+1)}) &\leq f(s_i - s_{i+1} + \sigma(s_j - s_i)t) - f(s_i - s_{i+1}) \\ &\quad + f(s_{j-1} - s_j + \sigma(s_j - s_i)t) - f(s_{j-1} - s_j) \\ &\leq \int_{s_i - s_{i+1}}^{s_i - s_{i+1} + \sigma(s_j - s_i)t} \frac{d\tau}{\sigma(-\tau)} \\ &\quad + \int_{s_{j-1} - s_j}^{s_{j-1} - s_j + \sigma(s_j - s_i)t} \frac{d\tau}{\sigma(-\tau)} \\ &\leq 2\int_z^{z+\sigma(-z)t} \frac{d\tau}{\sigma(-\tau)} \qquad (10) \\ &\leq \frac{2\sigma(-z)t}{\sigma(-z - \sigma(-z)t)} \\ &\leq \frac{2\sigma(-z)}{\sigma(-z - 2\sigma(-z))} t \end{aligned}$$

where (10) follows since $s_{j-1} \leq s_i$ and $s_{i+1} \geq s_j$. Lastly, we need to upper bound $\Phi(\mathbf{r}'; \pi_{\mathbf{r}'}) - \Phi(\mathbf{r}'; \pi_{\mathbf{r}})$. Note $\pi_{\mathbf{r}'}^{-1} \circ \pi_{\mathbf{r}}$ is at most two cycles along with a number of fixed points. The two sources of the cycles are when $s'_\ell < s'_i$ for $\ell < i$, i.e. the ith ranked player gained enough Elo to jump past the players numbered ℓ through $(i-1)$, and when $s'_j < s'_\ell$ for $j < \ell$, i.e. the jth player lost enough Elo to fall behind the players numbered $(j+1)$ through ℓ. In the first case, s'_i is moved to be placed just before s'_ℓ, so the relevant terms of $\Phi_e(\mathbf{r}'; \pi_{\mathbf{r}})$ are

$$f(s'_{\ell-1} - s'_\ell) + f(s'_{i-1} - s'_i) + f(s'_i - s'_{i+1})$$

and the relevant terms of $\Phi(\mathbf{r}'; \pi_{\mathbf{r}'})$ are

$$f(s'_{\ell-1} - s'_i) + f(s'_i - s'_\ell) + f(s'_{i-1} - s'_{i+1}).$$

Technically the above expressions assume $\ell > 1$, but since $f(s'_{\ell-1} - s'_i) \leq f(s'_{\ell-1} - s'_\ell)$, we can ignore the first term in those two expressions anyway and get a bound for all ℓ. Then since $s'_i > s'_{i-1}$, we have

$$f(s'_{i-1} - s'_i) + f(s'_i - s'_{i+1}) \geq f(s'_{i-1} - s'_{i+1}).$$

Consequently,

$$\Phi(\mathbf{r}';\pi_{\mathbf{r}'}) - \Phi(\mathbf{r}';\pi_{\mathbf{r}}) \le f(s'_i - s'_\ell).$$

But $s_\ell \ge s_i$ and $s'_i - s_i \le t$, so $f(s'_i - s'_\ell) \le f(t) \le t/\sigma(-1)$. The second case where the jth ranked player drops to ℓ follows by symmetry, and we're done. □

Lemma 5. *Let $r_{\max}$ denote the largest entry of $\mathbf{r}$. Then*

$$\Phi(\mathbf{r};\pi_{\mathbf{r}}) \ge \frac{1}{8}\frac{r_{\max}^3}{g^{-1}(r_{\max}^2/4)}$$

where g is given by (1).

Proof. For concision, put $s_i = r_{\pi_{\mathbf{r}}(i)}$ so $\Phi(\mathbf{r};\pi_{\mathbf{r}}) = \Phi(\mathbf{s};\mathrm{id})$. Then let $s_{i^*}, \cdots, s_{i^*+m-1}$ be a contiguous sub-sequence of $\mathbf{s}$ such that $s_{i^*} = s_{\max}$, each term is at least $s_{\max}/2$, and $s_{i^*+m} < s_{\max}/2$ (in other words, the sub-sequence is maximal). Since $s_n \le 0$, such a sub-sequence exists. Let $a_j := s_{i^*+j-1} - s_{i^*+j} \ge 0$. Then

$$\Phi(\mathbf{s};\mathrm{id}) \ge \|\mathbf{s}\|^2 + \sum_{i=i^*}^{i^*+m-1} f(s_i - s_{i+1}) = \|\mathbf{s}\|^2 + \sum_{j=1}^{m} f(a_j).$$

Note that we have the constraint $\sum_{j=1}^m a_j = s_{i^*} - s_{i^*+m} \ge s_{\max}/2$. By convexity of f, the quantity is minimized when the a_j are equalized. So,

$$\Phi(\mathbf{s};\mathrm{id}) \ge \|\mathbf{s}\|^2 + mf(s_{\max}/2m).$$

On the other hand, since $s_i \ge s_{\max}/2$ for $i \in \{i^*, \cdots, i^*+m-1\}$, we have $\|\mathbf{s}\| \ge m(s_{\max}/2)^2$. So we now seek to lower bound the quantity

$$\Phi(\mathbf{s};\mathrm{id}) \ge m(s_{\max}/2)^2 + mf(s_{\max}/2m). \tag{11}$$

To find the worse case m, substitute $m = 1/x$, differentiate with respect to x, and set the derivative to 0. The resulting equation is

$$(s_{\max}x/2)f'(s_{\max}x/2) = (s_{\max}/2)^2 + f(s_{\max}x/2).$$

The left-hand side is $(s_{\max}x/2)/\sigma(-s_{\max}x/2)$. Replace $z = s_{\max}x/2$ to obtain

$$\frac{z}{\sigma(-z)} - \int_0^z \frac{dt}{\sigma(-t)} = (s_{\max}/2)^2.$$

Applying integration by parts in reverse to the left-hand side shows it is in fact equal to $g(z)$. The worse-case value of m is therefore $(s_{\max}/2)/g^{-1}((s_{\max}/2)^2)$. Plugging this back into (11) gives

$$\Phi(\mathbf{s};\mathrm{id}) \ge \frac{s_{\max}^3}{8g^{-1}(s_{\max}^2/4)}.$$

□

These three lemmas can now be applied in sequence to give the main result of this section.

Theorem 6. *For each pot function σ, there exists some constant C such that $CR^3/g^{-1}(R^2/4)$ games must be played in order for a player to achieve a rating of at least R.*

Proof. Let p be any path starting at the origin and ending in T_R. Lemma 3 constructs an upset-free path p' with $\text{len}(p') \le C_1 \text{len}(p)$. Then, for any upset-free path p' starting at the origin and ending at $\mathbf{r}$, Lemma 4 gives $C_2 \text{len}(p') \ge \Phi(\mathbf{r}; \pi_{\mathbf{r}})$. Finally, $\Phi(\mathbf{r}; \pi_{\mathbf{r}}) \ge \frac{1}{8} R^3/g(R^2/4)$ for any $\mathbf{r} \in T_R$ by Lemma 5. All together, this shows

$$\text{len}(p) \ge \frac{1}{8C_1C_2} \frac{R^3}{g^{-1}(R^2/4)}.$$

Finally, we may let $C = \frac{1}{8C_1C_2}$ for the end result. ☐

Annoyingly, this expression is the 'dual' of the kind of result from the previous section. Previously, we asked how high a rating is achievable given a fixed number of games. Here, we answered how many games are required to achieve a particular rating. However, for many g we can rearrange the inequality, demonstrated by the corollary below. The corollary provides the form displayed in Table 1.

Corollary 7. *Let C be the pot-dependent constant from Theorem 6. If $g(x) \ge (ax)^2$ for some constant a for sufficiently large x, then the maximum rating achievable with k games and any number of players is at most*

$$C^{-1/3} k^{1/3} g^{-1} \left(k/(8aC)\right)^{1/3}.$$

for sufficiently large k.

Proof. Let R be the maximum rating achieved after k games.Then

$$\begin{aligned}
g\left(R/(2a)\right) \ge R^2/4 &\implies R/(2a) \ge g^{-1}\left(R^2/4\right) \\
&\implies \frac{1}{a} \frac{R^3/8}{g^{-1}\left(R^2/4\right)} \ge R^2/4 \\
&\implies g^{-1}\left(\frac{1}{8a} \frac{R^3}{g^{-1}\left(R^2/4\right)}\right) \ge g^{-1}\left(R^2/4\right) \\
&\implies \frac{R^3}{g^{-1}\left(R^2/4\right)} g^{-1}\left(\frac{1}{8a} \frac{R^3}{g^{-1}\left(R^2/4\right)}\right) \ge R^3 \\
&\implies \frac{R}{g^{-1}\left(R^2/4\right)^{1/3}} g^{-1}\left(\frac{1}{8a} \frac{R^3}{g^{-1}\left(R^2/4\right)}\right)^{1/3} \ge R
\end{aligned}$$

Define $h(x) = C^{-1/3} x^{1/3} g^{-1}(x/(8aC))^{1/3}$ so that the last inequality implies

$$h\left(\frac{CR^3}{g^{-1}(R^2/4)}\right) \ge R.$$

Then by Theorem 6, we have

$$\frac{CR^3}{g^{-1}(R^2/4)} \le k.$$

Applying h to both sides results in $R \le h(k)$ as desired. ☐

References

[1] Arpad E. Elo. *The Rating of Chess Players, Past and Present.* Arco Publishing, second edition, 1978.