CSE702, Spring 2024: Analyzing Cognitive Tendencies From Chess Data

Chess Ratings---Elementary musings based on an "Angry Statistician" post:

Suppose Player 1 has probability x of beating a generic opponent and Player 2 has probability y. Can we infer from x and y the probability p of Player 1 beating Player 2? We have some axioms:

1. $x = y \implies p = 0.5$.
2. $x = 0 \implies p = 0$ (maybe unless $y = 0$).
3. $x = 1 \implies p = 1$ (maybe unless $y = 1$).
4. $y = 0 \implies p = 1$ (maybe unless $x = 0$).
5. $y = 1 \implies p = 0$ (maybe unless $x = 1$).

It turns out we can derive a formula $p(x, y)$ with this behavior by dividing the "Player 1 odds ratio" $\frac{x}{1-x}$ by the Player 2 ratio $\frac{y}{1-y}$ to solve for the "direct confrontation odds ratio":

$$\frac{p}{1-p} = \frac{x(1-y)}{y(1-x)}$$

You can think of the odds ratio as the amount of money you need to bet to win \$1 when the payoff reflects the probability p. For instance, if $p = 0.75$ then the odds ratio is 3. If you bet \$1 and win the fair payoff is \$0.33... So you need to bet \$3 to win \$1 at this rate. Solving this for p gives $py(1-x) = x(1-y) - px(1-y)$, so $p(y - yx + x - xy) = x(1-y)$, so

$$p(x, y) = \frac{x(1-y)}{x + y - 2xy}.$$

[I verified in class that this satisfies the five axioms. See interesting question in notes at the end about the extent to which this formula may be unique according to the five axioms.]

We can actually derive this formula in a more elementary way that also takes into account the idea of an incremental struggle.

Consider the following possibilities for (1) a bowler in cricket or pitcher in baseball, versus (2) a ~~batsman~~ batter:

- Bowler/pitcher makes a good delivery: probability p_1.
- Bowler has poor length/pitcher "hangs" a curveball: $q_1 = (1 - p_1)$.
- Batter has good stroke, makes solid contact: p_2.
- Batter nicks or misses ball: q_2.

For sake of argument, we suppose that if both the delivery and the batter's stroke are good, the result is a dot-ball in cricket, or a foul ball in baseball, and the confrontation goes on. This is like both players making a good move at one game turn at chess. Or if the delivery and stroke are both bad, a mistimed hit (for no runs) or another foul ball may result. We get a *result* only when:

- Batter punishes a poor delivery: boundary or home run, probability $p_1 q_2$.
- Batter fails on a good delivery: wicket or strikeout, probability $p_2 q_1$.

The probability of the batter succeeding therefore is

$$\frac{p_1 q_2}{p_1 q_2 \; + \; p_2 q_1} \;=\; \frac{p_1(1-p_2)}{p_1(1-p_2) \; + \; p_2(1-p_1)} \;=\; \frac{p_1(1-p_2)}{p_1 \; + \; p_2 \; - \; 2\,p_1 p_2}\;.$$

This is the same formula as before with p_1 in place of x and p_2 in place of y.

Now we note a further twist. Divide both the numerator and denominator of the leftmost form of the equation by $p_1 q_2$. This gives the overall win probability of player 1 as:

$$\frac{1}{1 \; + \; \frac{p_2 q_1}{p_1 q_2}} \;=\; \frac{1}{1 \; + \; \left(\frac{p_2}{1-p_2}\right) / \left(\frac{p_1}{1-p_1}\right)}.$$

Now we have a ratio of two odds ratio feractions nestled inside another fraction. It looks weird, but now let's think more about the nature of an odds ratio $\frac{x}{1-x}$ as a mathematical function. It is always nonnegative and increases from zero to infinity as x goes from 0 to 1. This is the same range behavior as the exponential function e^M where M goes from $-\infty$ to $+\infty$, i.e., as a function of the whole real number line. In fact, the correspondence is exactly $M = \ln\left(\frac{x}{1-x}\right)$ which is the **logit** function, but let's not even think of that. Let's think of M abstractly as a measure of "mojo". A person who is more likely to lose than win ($x < 0.5$) has "negative mojo." An omnipotent player has infinite mojo, while a hopeless player has negative infinity mojo. If we substitute the "mojo" representations using M_1 and M_2 in place of the odds ratios for p_1 and p_2, we get:

$$\frac{1}{1 \; + \; e^{M_2} / e^{M_1}} \;=\; \frac{1}{1 \; + \; e^{(M_2 - M_1)}}.$$

The philosophical magic is this: We have converted the win probability of player 1 from a function of two variables representing the players separately into a function of *only one variable*: the "difference in mojo" between the players. This also means that the relation of winning probability to (difference in) "mojo" is the same across the scale.

(Note, incidentally, that this win probability is not meant to be the same as the "p_1" (or "x") we started with. The first time we derived the formula, x was the probability of winning against a "generic" opponent (or an average win rate over unspecified opponents), and y likewise for player 2 against general opposition; what we get is the probability p for player 1 against player 2 specifically. The second time, p_1 was a probability of personal success in isolation, which could involve skill factors apart from the quality of player 2's actions. And also by the way, we haven't yet said we are talking about *chess* or any other two-person strategy game. That chess has draws can be accommodated by the theory---we count "points expectation" instead of "win probability.")

When $M_2 > M_1$ the fraction is < 0.5, so player 1 is favored to win only when $M_1 > M_2$. What difference gives 75% win probability? Since $0.75 = \frac{1}{1+1/3}$ the answer is

$$M_1 - M_2 = \ln(3) = 1.0986...$$

Here is where I suspect that [Arpad Elo](), the "[Martian]()" who converted the notion of "mojo" into a statistically regulated rating system, indulged a little bit of "numerical voodoo" to make things look cleaner for the indigenous population he landed among. Since $10^x = e^{x \ln 10}$, we can change the base to be 10 (or any other number, but the humanoids have 10 fingers). Since we haven't specified what units "mojo" comes in, let us rewrite the player 1 success formula as

$$\frac{1}{1 + 10^{(M_2 - M_1)}}.$$

Now the answer we want is $M_1 - M_2 = \ln(3)/\ln(10) = 1.0986.../2.302585... = 0.47712...$ Hmmm...this is almost $1/2$. What happens if we plug in $M_2 - M_1 = -1/2$? We get

$$p = \frac{1}{1 + 1/\sqrt{10}} = \frac{1}{1 + 0.3162...} = \frac{1}{1.3162...} \sim = 0.7597...$$

Close enough to call this "75%"? This is so tempting, because if we want a nice round number D to mean the difference that gives "75%" probability, then our scaling factor can just be $2D$ in the denominator of the exponent, another nice round number. The US Chess Federation had already decided to call 200 points the width of a "class" under a rougher rating system devised by Kenneth Harkness in 1950, so Elo made $D = 200$ and the rating formula thus became the form it has today:

$$p = \frac{1}{1 + 10^{-(R_1 - R_2)/400}}.$$

László Mérő---who does not count as a "Martian" because he was born after WWII and stayed in Hungary---seized on the 75% advantage as a universal yardstick---a "Class Unit" of skill in any human endeavor. Being off by 0.97 percentage points may not seem a big deal, but consider this for humor:

Elo's fudge is the same as considering $\sqrt{10}$ to equal 3. The Hebrew Bible has passages that seem to equate $\pi = 3$. Well, $\pi = 3.14159...$, which is less of a stretch than $3.162...$ Thus *Elo* had greater chutzpah than *Elohim*.

The formula does make 200 into the "source" standard deviation of a player's rating. If we assume that all players are equally variable in their level of "mojo" at any given time, then the standard deviation of the difference $R_1 - R_2$ becomes $200\sqrt{2}$. Elo indulged two other fudges that help everything offset well enough, the first of which most data scientists allow generally:

- The slight unevenness between a logistic curve and the "probit curve", meaning the cumulant of the normal distribution, even after the famous "1.7" scaling factor is applied.
- The approximation $\sqrt{2} \;=\; 1.41421... \;\approx\; \frac{10}{7} \;=\; 1.42857...$ That at least took rather less "chutzpah"! Thus he represented $200\sqrt{2}$ as $2000/7$.

The closeness of the resulting nexus of Elo's logistic-based probabilities and their conformance to normal distribution is shown in my GLL blog article "Sliding-Scale Problems" (original source is François Labelle's "Elo Probability Win Calculator"), which is on the upcoming seminar menu:

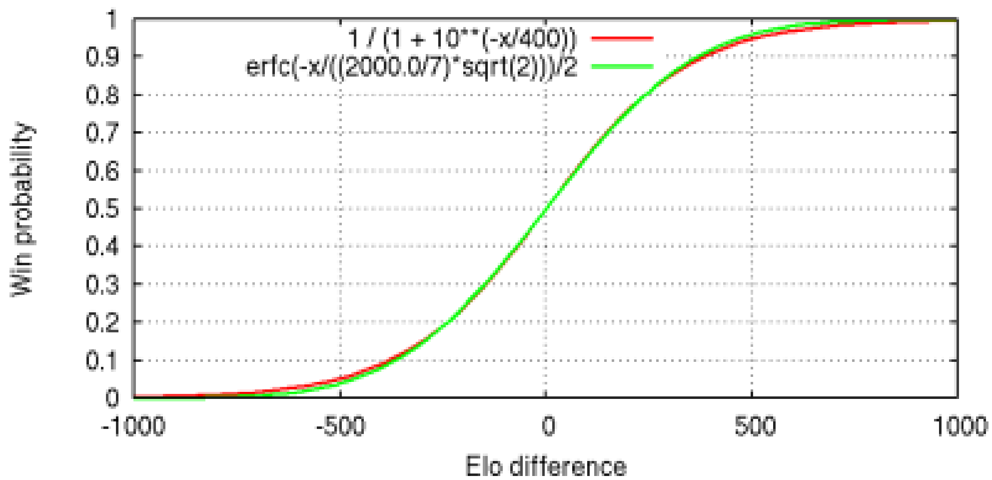

See also Nate Solon's article "How Elo Ratings Actually Work." That Elo didn't care about super-fine precision is witnessed by his famous summary of the whole shebang:

"The process of rating players can be compared to the measurement of the position of a cork bobbing up and down on the surface of agitated water with a yard stick tied to a rope which is swaying in the wind." (quote source).

But this is in how his system is *applied*. Speaking as a mathematical Platonist, I find the logistic formula to be **salient**---and thus "divinely ordained" as a matter of *theory*. This extends to my belief that quantities that are strong "telltales" of a player's "mojo" should be linear in it across the entire scale, full-stop. They should certainly not be "kinked" as in these two diagrams with the same data points (representing "T3" match to Stockfish 16 with 0.5 inferiority cutoff):

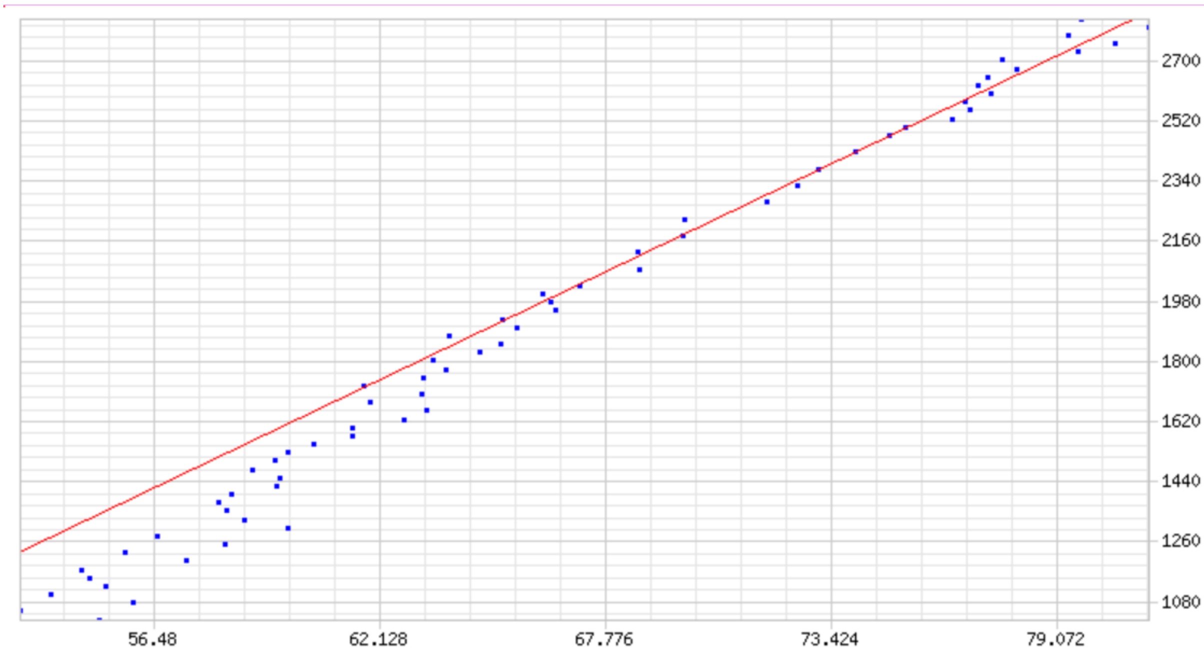

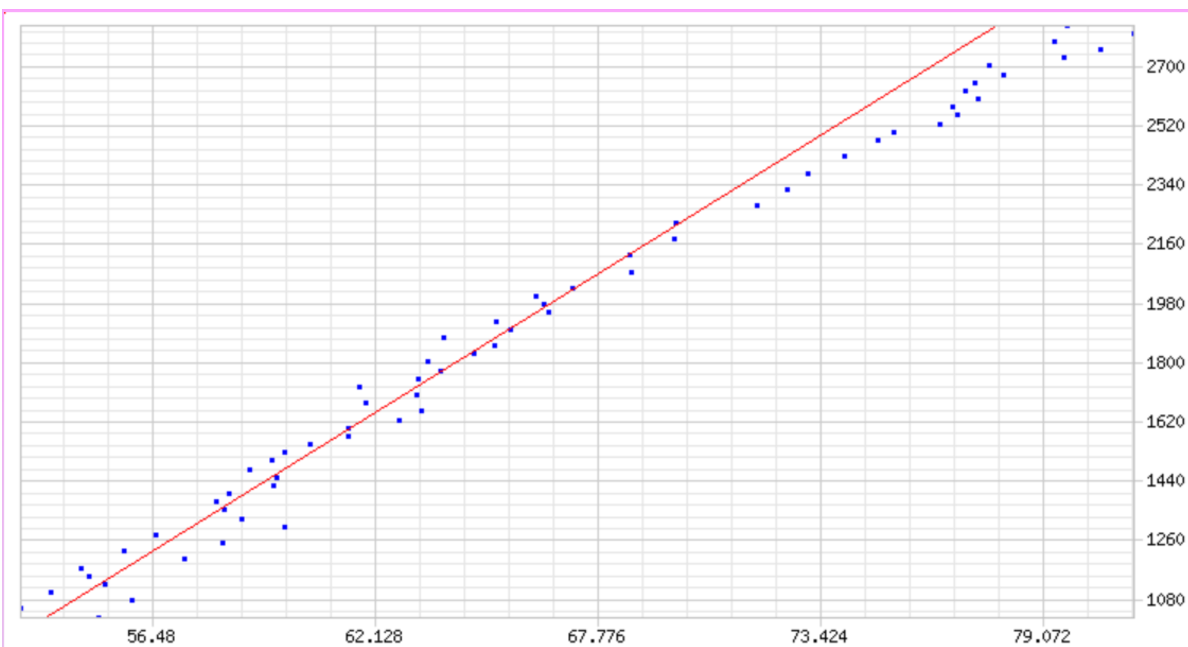

Footnote:

An interesting question is whether this formula is unique for **any ratio of** two possibly-infinite power series in x and y. Note that power series in just x alone encompass exponentiation and logarithms and all trig functions. So a two-dimensional power series in both x and y, and a ratio of the same, with arbitrary real coefficients, is quite a general mathematical function. To get a start on this idea, write the ratio in general form as

$$\frac{\sum_{i,j\geq 0} a_{ij}x^iy^j}{\sum_{i,j\geq 0} b_{ij}x^iy^j} \;=\; \frac{a_{00} \;+\; \sum_{i\geq 1} a_{i0}x^i \;+\; \sum_{j\geq 1} a_{0j}y^j \;+\; \sum_{i,j\geq 1} a_{ij}x^iy^j}{b_{00} \;+\; \sum_{i\geq 1} b_{i0}x^i \;+\; \sum_{j\geq 1} b_{0j}y^j \;+\; \sum_{i,j\geq 1} b_{ij}x^iy^j}.$$

Then Axiom 2 says that when $x = 0$, the entire numerator must vanish whatever y is (except that the case $y = 0$ is allowed to be indeterminate). Therefore all the coefficients a_0j with $j \geq 1$ must be identically zero, else y could make it vary. And we must have the constant term $a_{00} = 0$ too. Once you whittle down the terms this way with axioms 2--5, axiom 1 will step in to say that for each n, the sum of b_{ij} over $i + j = n$ must be exactly twice the sum of a_{ij} over $i + j = n$. Maybe it might follow that those sums must be identically zero for $n \geq 3$. Well, you could also suppose the terms with $i + j \geq 3$ are absent to begin with---i.e., that $p(x, y)$ is a ratio of quandratic polynomials. Then must the above formula be the only possibility?