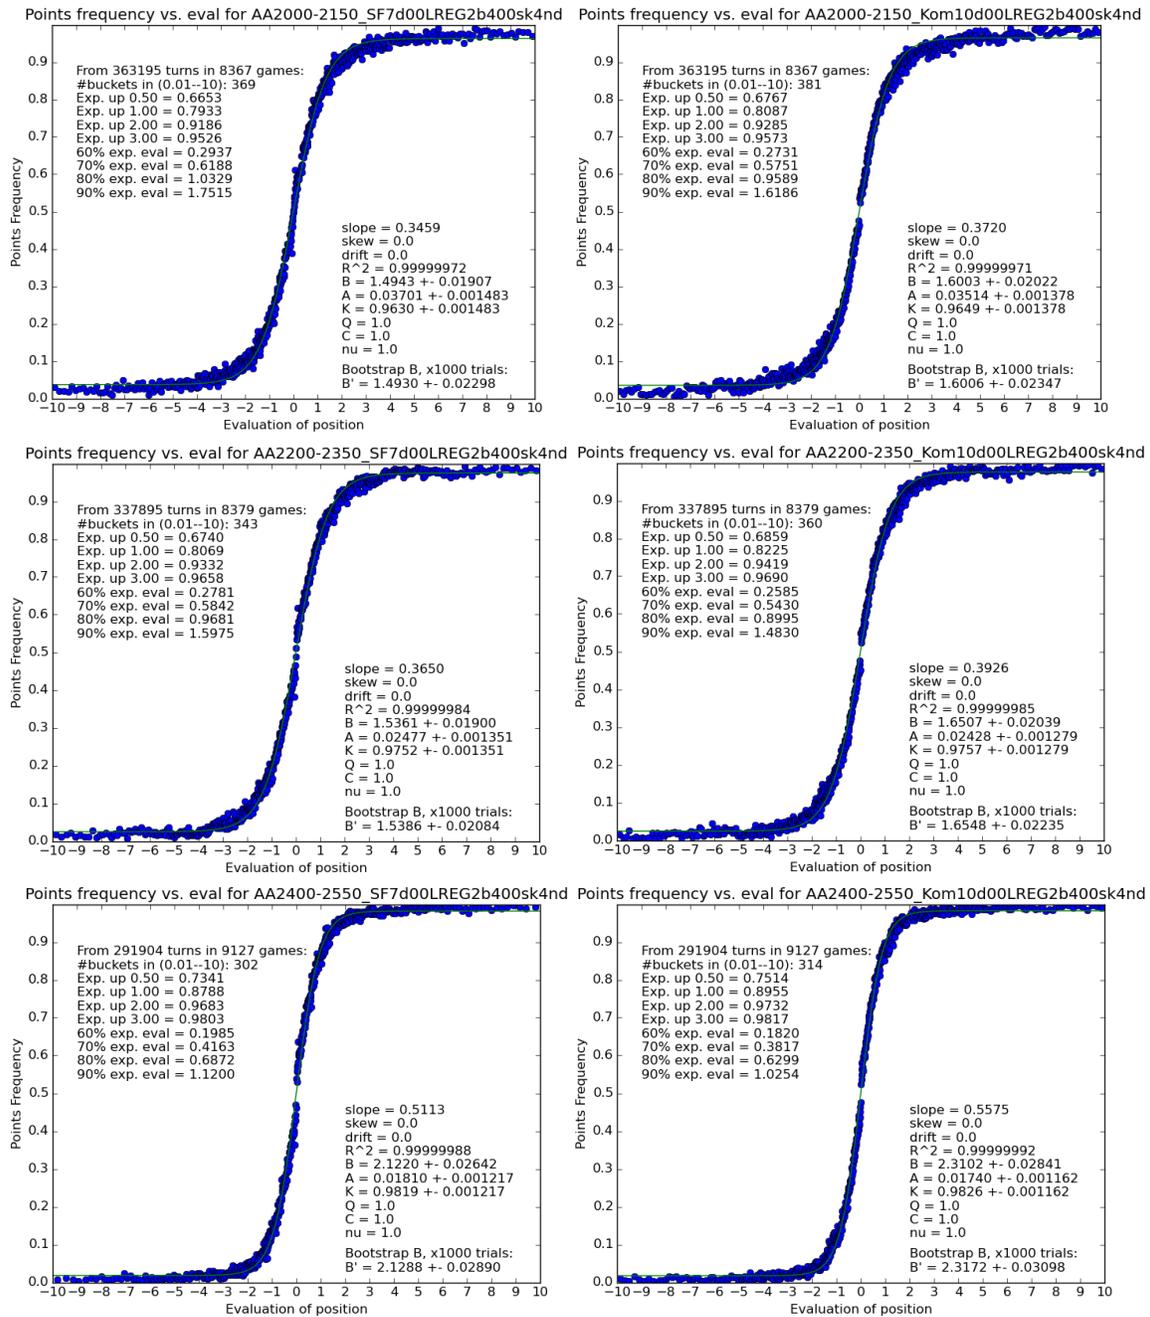
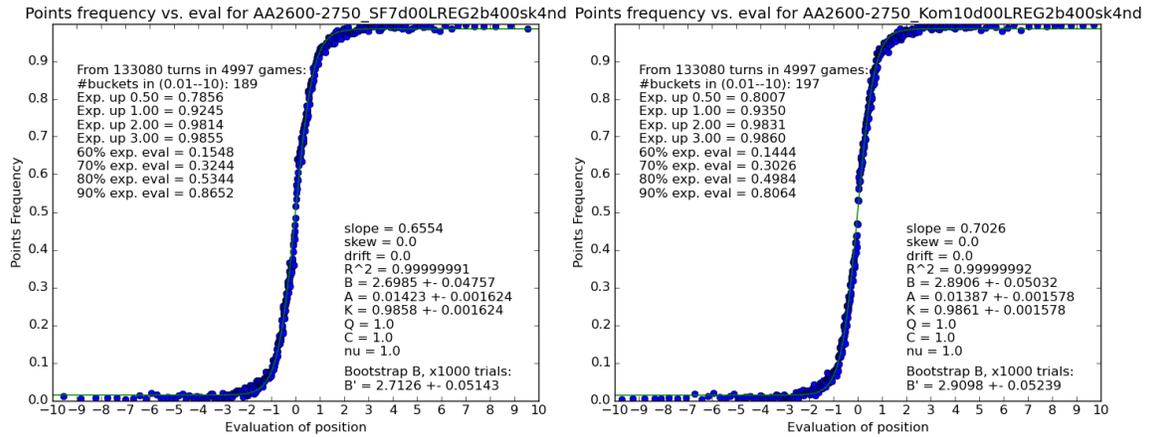


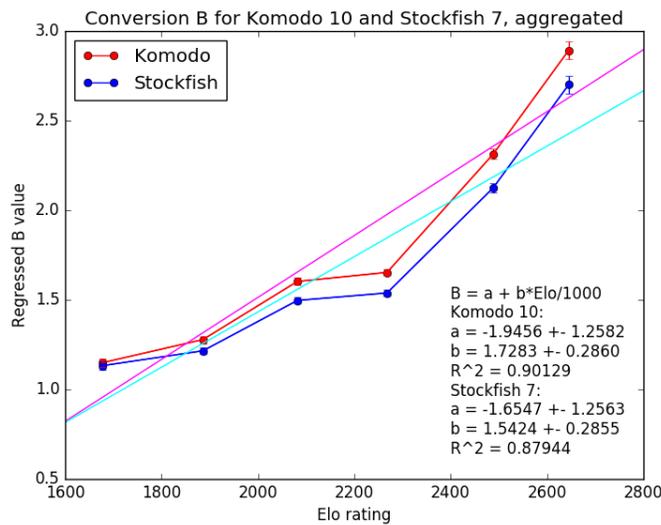
Runs on Decisive Games Only

The logistic relation holds even when drawn games are removed from the dataset. Here are side-by-side comparisons of Stockfish 7 (left) and Komodo 10 (right) for Elo 2000–2150 through 2600–2750; the “nd” at the end stands for “no draws”:

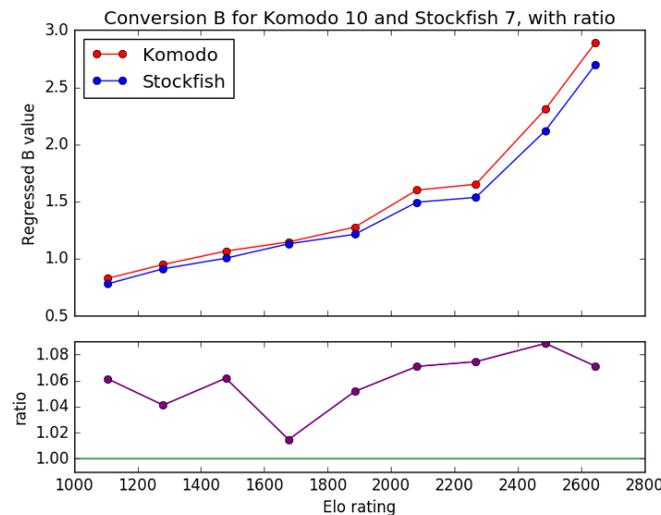


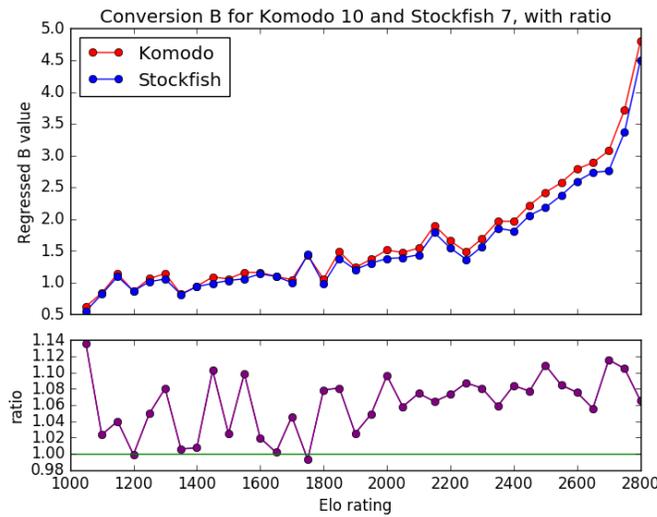


The skewness of the data point for 2200–2350 (plotted at 2282) is if anything accentuated.



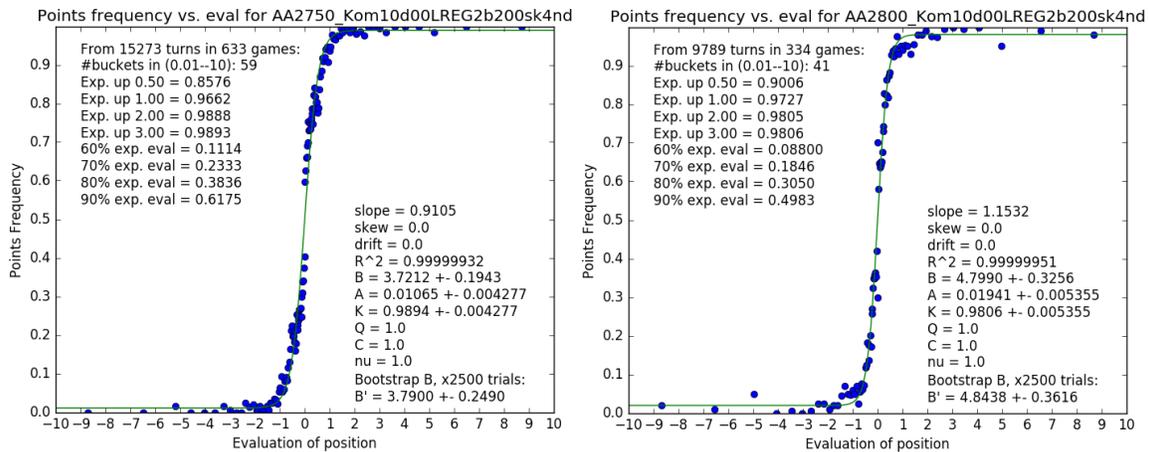
The point for 2600–2750 does not go down. Here are the plots of B and its Komodo/Stockfish ratio over all the groupings, then for each Elo level individually:





The first plot briefly gives hope that the ratio 1.06 is workable for the whole range, downplaying the outlier for 1600–1750. The second however shows that it zooms over 1.10 for several levels throughout the range and wobbles wildly overall.

Finally here are the individual plots for the 2750 and 2800 with their high spikes, shown for Komodo:



Note the drastically lower number of turns—because a higher percentage of games at these levels are drawn. The curve is showing a high spike for slight advantages *conditioned on the game being decisive*. That the results still cling so closely to a pure logistic curve is remarkable.